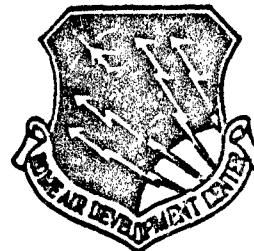


AD-A154 033

RADC-TR-85-55

Final Technical Report

March 1985



***LAN INTEROPERABILITY STUDY OF
PROTOCOLS NEEDED FOR DISTRIBUTED
COMMAND AND CONTROL***

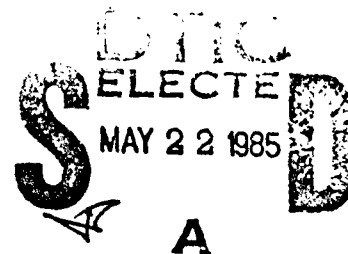
Harris Corporation

Walter L. Elden, Anita L. Miller, Stephen L. Morgan
and Barbara A. Romanzo

APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED

DTIC FILE COPY

20030116012



**ROME AIR DEVELOPMENT CENTER
Air Force Systems Command
Griffiss Air Force Base, NY 13441-5700**

This report has been reviewed by the RADC Public Affairs Office (PA) and is releasable to the National Technical Information Service (NTIS). At NTIS it will be releasable to the general public, including foreign nations.

RADC-TR-85-55 has been reviewed and is approved for publication.

APPROVED:

Thomas F. Lawrence
THOMAS F. LAWRENCE
Project Engineer

APPROVED:

Raymond P. Urtz, Jr.
RAYMOND P. URTZ, JR.
Technical Director
Command and Control Division

FOR THE COMMANDER:

John A. Ritz
JOHN A. RITZ
Acting Chief, Plans Office

If your address has changed or if you wish to be removed from the RADC mailing list, or if the addressee is no longer employed by your organization, please notify RADC (COTD) Griffiss AFB NY 13441-5700. This will assist us in maintaining a current mailing list.

Do not return copies of this report unless contractual obligations or notices on a specific document requires that it be returned.

REPRODUCTION QUALITY NOTICE

This document is the best quality available. The copy furnished to DTIC contained pages that may have the following quality problems:

- Pages smaller or larger than normal.
- Pages with background color or light colored printing.
- Pages with small type or poor printing; and or
- Pages with continuous tone material or color photographs.

Due to various output media available these conditions may or may not cause poor legibility in the microfiche or hardcopy output you receive.

☐

If this block is checked, the copy furnished to DTIC contained pages with color printing, that when reproduced in Black and White, may change detail of the original copy.

UNCLASSIFIED
SECURITY CLASSIFICATION OF THIS PAGE

REPORT DOCUMENTATION PAGE				
1a. REPORT SECURITY CLASSIFICATION UNCLASSIFIED		1b. RESTRICTIVE MARKINGS N/A		
2a. SECURITY CLASSIFICATION AUTHORITY N/A		3. DISTRIBUTION/AVAILABILITY OF REPORT Approved for public release; distribution unlimited.		
2b. DECLASSIFICATION/DOWNGRADING SCHEDULE N/A				
4. PERFORMING ORGANIZATION REPORT NUMBER(S) N/A		5. MONITORING ORGANIZATION REPORT NUMBER(S) RADC-TR-85-55		
6a. NAME OF PERFORMING ORGANIZATION Harris Corporation	6b. OFFICE SYMBOL (If applicable)	7a. NAME OF MONITORING ORGANIZATION Rome Air Development Center (COTD)		
6c. ADDRESS (City, State and ZIP Code) Government Information Systems Division P.O. Box 98000 Melbourne FL 32902		7b. ADDRESS (City, State and ZIP Code) Griffiss AFB NY 13441-5700		
8a. NAME OF FUNDING/SPONSORING ORGANIZATION Rome Air Development Center	8b. OFFICE SYMBOL (If applicable) COTD	9. PROCUREMENT INSTRUMENT IDENTIFICATION NUMBER F30602-83-C-0108		
8c. ADDRESS (City, State and ZIP Code) Griffiss AFB NY 13441-5700		10. SOURCE OF FUNDING NOS.		
		PROGRAM ELEMENT NO. 62702F	PROJECT NO. 5581	TASK NO. 21
				WORK UNIT NO. 55
11. TITLE (Include Security Classification) LAN INTEROPERABILITY STUDY OF PROTOCOLS NEEDED FOR DISTRIBUTED COMMAND AND CONTROL				
12. PERSONAL AUTHOR(S) Walter L. Elden, Anita L. Miller, Stephen L. Morgan, Barbara A. Romanzo				
13a. TYPE OF REPORT Final	13b. TIME COVERED FROM Jun 83 TO Jul 84	14. DATE OF REPORT (Yr., Mo., Day) March 1985	15. PAGE COUNT 340	
16. SUPPLEMENTARY NOTATION N/A				
17. COSATI CODES			18. SUBJECT TERMS (Continue on reverse if necessary and identify by block number)	
FIELD	GROUP	SUB. GR.		
09	02		Local Area Network Protocols, Distributed Systems, Computer Networks; e d	
19. ABSTRACT (Continue on reverse if necessary and identify by block number) The study examined distributed processing requirements for strategic and tactical C ³ I systems, reviewed the characteristics and architectural issues for distributed processing global operating systems, compared the DoD and ISO networking protocol architecture models, the protocols for LAN's developed by the IEEE and ANSI, reviewed and conducted performance evaluation of Ethernet, DoD's Internet Protocol and Transmission Control Protocol and reported characteristics of CSMA/CD, Token Bus and Token Ring LAN's, reviewed three alternatives to using TCP for an intra-LAN protocol and examined the methods for employing gateway elements to interconnect LAN-based system elements. A comprehensive discussion of the results is given followed by a set of concise conclusions. Ten recommendations are given, providing a roadmap to guide the Air Force in developing C ³ I systems and LAN-based protocols. Three major areas are identified where future work is needed. A set of protocols and design approaches for internetworking is contained in a set of appendices.				
20. DISTRIBUTION/AVAILABILITY OF ABSTRACT UNCLASSIFIED/UNLIMITED <input checked="" type="checkbox"/> SAME AS RPT. <input type="checkbox"/> DTIC USERS <input type="checkbox"/>			21. ABSTRACT SECURITY CLASSIFICATION UNCLASSIFIED	
22a. NAME OF RESPONSIBLE INDIVIDUAL Thomas F. Lawrence			22b. TELEPHONE NUMBER (Include Area Code) (315) 350-2158	22c. OFFICE SYMBOL RADC (COTD)

DD FORM 1473, 83 APR

EDITION OF 1 JAN 73 IS OBSOLETE.

UNCLASSIFIED
SECURITY CLASSIFICATION OF THIS PAGE

Keywords included:

PREFACE

This is the final technical report for a LAN Interoperability Study conducted for the Rome Air Development Center of the Air Force Systems Command under contract F30602-83-C-0108. The study investigated issues associated with LAN-based systems protocols needed for building C³I distributed information processing systems. The study spanned over 12 months and devoted 3000 professional hours of study effort.

The report sets forth an introduction to the compiled material, reviews some background of the problem, sets forth what the study objectives were and discusses the approach taken. A comprehensive discussion of the results is given followed by a set of concise conclusions. Ten recommendations are given, providing a roadmap to guide the Air Force in developing C³I systems and LAN-based protocols. Three major areas are identified where future work is needed. A set of protocols and design approaches for internetworking is contained in a set of appendices.

The Principal Investigator and Study Manager for the study was Walter L. Elden. The leader of the simulation and modeling effort was Anita L. Miller. She was assisted by Susan West and Sheila Kasprzak. Stephen Morgan contributed to the study of local and network operating systems. Barbara Romanzo reviewed the DOD's high level protocols and TCP. Dr. Peter Knoke contributed in reviewing the NOS and DOS forms of global operating systems. Dave Carson provided consulting advice on distributed operating systems. William Windham contributed to the early simulation and modeling work. Dr. Larry King contributed the material dealing with multilevel security issues.

Mr. Thomas F. Lawrence, of the Distributed Systems Section, C² Systems Technology Branch, was the Contracting Officer's Technical Representative during the course of the study contract.

0405b/LAN

1



Accession For	
NTIS GRA&I	<input checked="" type="checkbox"/>
DTIC TAB	<input type="checkbox"/>
Unannounced	<input type="checkbox"/>
Justification	
By	
Distribution/	
Availability Codes	
Dist	Avail and/or Special
A1	

SUMMARY

The LAN Interoperability Study Final Report presents the results of a year long, 3000 professional study hours, investigation of issues associated with Local Area Network (LAN) based protocols needed for C³I distributed processing. A roadmap, comprising ten recommendations, three major areas of future study and eight appendices, sets forth a plan aimed at focusing the resources of the Air Force to developing distributed processing systems for the 1990's C³I programs.

Distributed, secure, survivable information systems are needed for the strategic/tactical battlefields to enable Air Force personnel to maintain control over forces, provide intelligence about enemy intentions and capabilities, warn of attacks or hostile actions, help conserve resources and aid with countless other tasks. The key to achieving these capabilities lies in applying distributed systems technologies that are combinable to create an integrated, system-wide command, control, communications and intelligence capability (C³I).

LAN-based networking protocols will be needed for distributed information processing systems. To date, which ones and how they would differ from Wide Area Network (WAN) protocols has not been well understood. It is recognized, though, that the applications environment in which these protocols will be employed will be one exhibiting a high degree of heterogeneity. The component elements (i.e., computers, peripherals, terminals), the end system as a whole and the LAN's and WAN's will probably differ in their architectures, interfaces and protocols.

The study examined distributed processing requirements for strategic and tactical C³I systems, reviewed the characteristics and architectural issues for distributed processing global operating systems, compared the DOD and ISO networking protocol architecture models, the protocols for LAN's developed by the IEEE and ANSI, reviewed and conducted performance evaluation of Ethernet, DOD's Internet Protocol and Transmission Control Protocol and reported characteristics of CSMA/CD, Token Bus and Token Ring LAN's, reviewed three alternatives to using TCP for an intra-LAN protocol and examined the methods for employing gateway elements to interconnect LAN-based system elements.

The following are the main conclusions reached by the study:

1. C³ must survive and offer sufficient flexibility to identify, reconstitute, and employ surviving assets for trans- and post-attack command and control.
2. For the 1990's, tactical C² systems will need to be interoperable across the military services and their many operational and support systems.
3. Currently, use of DOD's TCP and IP protocols appears justified to ensure interoperability when interworking through wide area networks. Intra-LAN use needs further consideration.
4. The International Standards Organization work, in developing the OSI/RM and its suite of protocols, was seen as the key indicator of where industry was going with protocols.
5. The Air Force's Master Plan for the 1990's (TAFIIS) calls for a dispersed, distributed, survivable C³ system. The plan's concept calls for configuration in a distributed, modular architecture with sharing of information seen as the key to surveillance and intelligence effectiveness.
6. C³ needs for the tactical Army's Air/Land Battle 2000 system are very much the same as needed for the Tactical Air Force.
7. Heterogeneous processors need to be interoperable over networks, employing high level protocols and packet-switching to achieve distributed processing for C³I.
8. A global Network Operating System (NOS) or Distributed Operating System (DOS) form of high level operating system is required to manage the distributed system resources. The NOS form builds on top of the original heterogeneous Constituent or Local Operating System, while the DOS form replaces it with a uniform homogeneous global one.
9. The National Software Works implemented the NOS form on top of the ARPANET Wide Area Network.
10. The CRONUS DOS is an NOS form being developed to work on top of Local Area Networks.
11. A Generic Network Operating System, called GNOS, was defined and provides a general reference model for developing protocols for C³I distributed processing systems.

12. The DOD network reference model provides the more basic networking utility services/protocols, whereas the ISO set is more comprehensive and is directed more formally to apply object-based design for future distributed processing applications.
13. The IEEE 802 Project has developed the currently leading industry protocols for LAN's, operating at 10 Mb/s.
14. ANSI X3T95 is developing a fiber-optic-based Token Ring LAN operating at 100 Mb/s.
15. The Air Force is developing the Flexible Intraconnect LAN (FILAN) to operate at 180 Mp/s and be used in C³I systems.
16. Based on analysis, the CSMA/CD offers an acceptable contention scheme for light to moderate loads while the Token Access Methods do for the deterministic approach, when moderate to heavy loads and controlled delay are criteria.
17. Simulation results indicated that throughputs up to 5 Mb/s were obtained for a single TCP connection on Ethernet (10 Mp/s) with no collision occurring and that protocol processing times contributed more to performance than did protocol overhead at the cable level.
18. Three alternatives to TCP for intra-LAN use were 1) extended backplane, 2) subset of TCP, and 3) a null layer.
19. A family of gateway elements used with open protocols was identified to achieve LAN-based systems interoperability.

Ten recommendations are made which provide the Air Force a roadmap for developing C³I distributed systems for the 1990's. Areas of future study are given and key protocols and internetworking/interoperability design approaches are presented in a set of appendices. The main thrusts of the recommendations are as follows:

1. Forming a joint Air Force-Industry C³I protocol development effort.
2. Developing a layered reference model for C³I applications and systems.
3. Employing the Generic Network Operating System (GNOS) to guide development of networking protocols for C³I.
4. Employment of gateway elements and GNOS protocols to achieve LAN-based systems interoperability.
5. Need for different intra-LAN and inter-LAN protocols for the underlying transport services.

6. Use of multiple LAN's to meet application requirements.
7. Continuing research into protocol design, validation and formal verification methods.
8. Need for a design practices handbook for quantifying LAN performance characteristics.
9. Need to evaluate performance of high level, multiple LAN's and WAN protocols through simulation and modeling.
10. Need to study further the areas of:
 - a. Distributed Systems Design Issues
 - b. Integrated OSI for Distributed Information Processing
 - c. Multilevel Secure Distributed Operating System Issues

TABLE OF CONTENTS

<u>Paragraph</u>	<u>Title</u>	<u>Page</u>
1.0	INTRODUCTION	1-1
2.0	BACKGROUND	2-1
3.0	STUDY OBJECTIVES	3-1
3.1	Initial Study Investigation Objectives	3-1
3.2	Revised Study Investigation Objectives	3-2
4.0	STUDY APPROACH	4-1
4.1	Assessment of DOD, Air Force and Industry Directions	4-1
4.1.1	Department of Defense Direction	4-1
4.1.2	Air Force Direction	4-2
4.1.3	Industry Directions	4-3
4.2	Tactical Command and Control System Models and Requirements	4-4
4.3	Distributed Information Processing	4-4
4.4	Operating Systems (Local and Global)	4-5
4.5	High Level Protocols Required for GNOS	4-6
4.6	Networking Architecture and Protocol Suites (ARPANET/DOD and ISO/OSI)	4-6
4.7	Performance Characteristics of Local Area Networks	4-6
4.8	Simulation and Modeling Performance Evaluation of TCP/ICP/LAN Protocols	4-7
4.9	Generic Gateways for Interoperability	4-9
4.10	Reporting of Interim Study Results	4-9
5.0	STUDY RESULTS	5-1
5.1	Department of Defense Study Findings	5-1
5.1.1	Strategic Distributed and Survivable Command and Control	5-1
5.1.2	Interoperability of 1990's Tactical Command and Control Systems	5-4
5.1.3	Tactical Command Control and Communications for the U.S. Army ..	5-8
5.1.4	Department of Defense Protocols Policy	5-10
5.2	Air Force Study Findings	5-14
5.2.1	Air Force Protocols Policy	5-14
5.2.2	Air Force LAN Architecture	5-16
5.2.3	Air Force Local Area Network Systems Program Office (AFLAN SPO)	5-20
5.3	Industry Study Findings	5-20
5.3.1	Open Systems Interconnection	5-20
5.4	Command, Control and Communications for Tactical Air Control System	5-24
5.4.1	Tactical Command and Control	5-24
5.4.2	Discussion	5-24
5.4.3	Tactical Air Control Missions and Architecture	5-25
5.5	Command Control and Communications for the Tactical Army System	5-28
5.5.1	General Network Architecture	5-29

TABLE OF CONTENTS (Continued)

<u>Paragraph</u>	<u>Title</u>	<u>Page</u>
5.5.2	Subsystem Architecture	5-31
5.6	Distributed Information Processing for Command, Control and Communications	5-32
5.6.1	General	5-32
5.6.2	Requirements for C ³ Distributed Information Processing	5-32
5.6.3	Distributed Processing Architectural Criteria	5-44
5.7	Operating Systems for C ³ Distributed Information Processing	5-46
5.7.1	An Operating System	5-46
5.8	National Software Works (NSW) Network Operating System	5-51
5.9	The Cronus Distributed Operating System	5-55
5.9.1	Introduction	5-55
5.9.2	Cronus, A Distributed Operating System	5-55
5.9.3	The Cronus DOS Functions	5-56
5.9.4	DOS Provides Essential Services System-Wide	5-59
5.9.5	Cronus Logical Model Architecture	5-60
5.9.6	Cronus Cluster Physical Model	5-62
5.9.7	DOS Generic Computer Elements	5-62
5.9.8	Interprocess Communication (IPC)	5-63
5.9.9	The Communication Subsystem	5-65
5.9.10	Candidate Protocols for Cronus	5-66
5.10	Generic Network Operating System (GNOS)	5-66
5.10.1	Introduction	5-66
5.10.2	Objectives for GNOS	5-67
5.10.3	GNOS Architecture	5-67
5.10.4	GNOS Subsystems	5-69
5.10.5	GNOS Services	5-69
5.10.6	GNOS Functions	5-70
5.10.7	Protocols Needed to Support GNOS	5-70
5.11	Comparison of DOD and ISO Networking Protocol Reference Models	5-73
5.11.1	Layered Architecture	5-74
5.11.2	DOD Versus ISO Models	5-78
5.12	DOD Networking Reference Model	5-83
5.13	ISO and IEEE 802 Networking Reference Models and Protocols	5-89
5.13.1	Introduction	5-89
5.13.2	ISO's Open System Interconnection (OSI) Architecture and Protocol Suite	5-92
5.13.3	Seven OSI Protocol Layers	5-93
5.13.4	IEEE Project 802 Protocols Suite for Local Area Networks	5-97
5.14	ANSI 100 Mb/s Token Ring LAN Standard	5-108
5.14.1	Introduction	5-108
5.14.2	Discussion of FDDI	5-108
5.15	Air Force Flexible Intraconnect Local Area Network (FILAN)	5-111
5.15.1	Introduction	5-111
5.15.2	System Overview	5-111
5.15.3	Air Force Workshop on Flexible Intraconnect Local Area Network (FILAN)	5-113
5.15.4	FILAN In a Multi-Media Environment	5-115
5.16	Effects of LAN Protocol Characteristics	5-115
5.16.1	Topological Effects [34]	5-115
5.16.2	Transmission [34]	5-116

TABLE OF CONTENTS (Continued)

<u>Paragraph</u>	<u>Title</u>	<u>Page</u>
5.16.3	Traffic Effects [34]	5-117
5.16.4	Performance Effects of LAN Access Methods	5-119
5.16.5	Comparing Three IEEE 802 LAN Systems Performance	5-121
5.17	Evaluation of TCP/IP in a Local Network Environment	5-132
5.17.1	Performance Results	5-132
5.17.2	Simulation/Modeling of TCP/IP/Ethernet LAN	5-133
5.18	TCP Alternatives for Intr-LAN Use	5-152
5.18.1	The Local Network Transmission Control Protocol (LNTCP) Alternative	5-152
5.18.2	An Extended Backplane Approach to LAN Interprocess Communications	5-155
5.18.3	Protocol Functional Approach to LAN	5-156
5.19	Generic Gateways for LAN Interoperability	5-157
5.19.1	Introduction	5-157
5.19.2	Objectives	5-158
5.19.3	Open Systems Interconnection	5-160
5.19.4	LAN Architecture and Protocols	5-160
5.19.5	Connectionless and Connection-Oriented Services [61]	5-161
5.19.6	Principles of Interconnection	5-162
5.19.7	Gateways for Interoperability	5-164
6.0	CONCLUSIONS	6-1
6.1	General	6-1
6.2	Strategic Command, Control and Communications	6-1
6.3	Tactical Command and Control-Interoperability and Survivability	6-1
6.4	Protocol Standards for Military Use	6-2
6.5	Protocol Standards for Industry Use	6-2
6.6	Tactical Air Control Command, Control and Communications	6-3
6.7	C ³ for the Tactical Army System	6-4
6.8	Distributed Information Processing for C ³	6-4
6.9	Operating Systems C ³ Distributed Information Processing	6-5
6.10	National Software Works Network Operating Systems	6-5
6.11	Cronus DOS	6-6
6.12	Generic Network Operating System (GNOS)	6-6
6.13	Comparison of DOD and ISO Networking Protocol Reference Models	6-7
6.14	ANSI 100-Mb/s LAN	6-8
6.15	Air Force Flexible Intraconnect LAN (FILAN)	6-9
6.16	LAN Protocol Characteristics and Effects	6-9
6.17	Evaluations of TCP/IP in a LAN Environment	6-11
6.18	TCP Alternatives for Intra-LAN Use	6-12
6.19	Generic Gateways for LAN Interoperability	6-13
7.0	RECOMMENDATIONS	7-1
8.0	AREAS OF FUTURE WORK	8-1
8.1	General	8-1
8.2	Distributed System Design Issues	8-1

TABLE OF CONTENTS (Continued)

<u>Paragraph</u>	<u>Title</u>	<u>Page</u>
8.2.1	Introduction	8-1
8.2.2	Overall Architectural Model	8-2
8.2.3	Specific Design Issues	8-4
8.2.4	Conclusions	8-9
8.3	Integrated Open Systems Interconnection for Distributed Information Processing	8-10
8.3.1	ANSI/OSI Networking Committee Reorganization	8-10
8.3.2	Integrated OSI for Distributed Information Processing	8-13
8.3.3	Issues Associated with an Integrated OSI for Distributed Information Processing	8-14
8.4	Multilevel Secure Distributed Operating System	8-15
8.4.1	Introduction	8-15
8.4.2	Cronus DOS Baseline	8-16
9.0	REFERENCES	9-1

APPENDICES

A	Evaluation of DOD Higher Layer Protocols	A-1
B	Some Improvements to the DOD Higher Layer Protocols	B-1
C	Transmission Control Protocol (TCP) Usage for LANs	C-1
D	Protocols Identified for the Generic Network Operating System (GNOS) Reference Model	D-1
E	Networking and System Resource Management Identified Protocols	E-1
F	Remote Data Base Access Protocol	F-1
G	Generic Gateways for LAN Interoperability	G-1
H	Multimedia LAN (MMLAN) Internetworking	H-1

LIST OF ILLUSTRATIONS

<u>Figure</u>	<u>Title</u>	<u>Page</u>
2.0	LAN-Based Distributed Command and Control Network Model	
	Problem Investigated	2-2
3.1-1	Intra-LAN (MINIDOS) Cluster of Heterogeneous Host Machines (Candidate Baseline C ² System Model)	3-3
3.1-2	Inter-LAN (MAXIDOS) Cluster-to-Cluster of Heterogeneous Host Machines (Candidate Baseline C ² System Model)	3-3
3.1-3	Layered Architecture for Communications Network Protocols ..	3-3
3.1-4	Protocol Hierarchy Employing TCP/IP	3-3
5.1.1	Partitioned Islands of Surviving C ³ Resources	5-2
5.1.1.1-1	Distributed Processing Layered Architecture for Survivable C ³	5-3
5.1.2-1	Automated Battlefield System Planned for the 1990's	5-5
5.1.2-2	Proposed Common User Architecture for Tactical Information Exchange System	5-6
5.1.2-3	DOD Internet Protocol Hierarchy	5-7
5.1.4.1	DOD Internet Protocol Hierarchy	5-13
5.3.1-1	The Seven Layer OSI Architecture	5-21
5.3.1-2	Open Systems Interconnections Between Closed Systems	5-22
5.4.3-1	Air Control and Air Surveillance Systems Operational Organization	5-26
5.4.3-2	Distributed Air Control and Netted Air Surveillance System Concept	5-27
5.4.3-3	TACS Modular Operations Concept Example Deployment	5-28
5.4.3-4	Command and Control Module (CCM) Shelter Layout Concept	5-29
5.4.3-5	Maximum Modular Operations Concept (MOC) Configuration	5-31
5.6.2.2	Command Center Protocol Architecture	5-35
5.6.2.2-1	Logical View of the TAFIIS Architecture	5-38
5.6.2.2-2	Physical View of the TAFIIS Data Processing Architecture ...	5-39
5.6.2.2-3	Conceptual User Interface Architecture	5-40
5.8	National Software Worker (NSW) Form of Network Operating System	5-52
5.9.2.2	The Local Cronus Cluster Configuration (Physical Model)	5-57
5.9.2.3	The Cronus Inter-Cluster Environment (Physical Model)	5-58
5.9.5-1	Cronus DOS Architecture (Logical View)	5-60
5.9.5-2	Distributed System Environment	5-61
5.9.8	Cronus Interprocess Communication	5-63
5.10.3	Generic Network Operating System (GNOS)	5-68
5.11.2	DOD and ISO Reference Models	5-78
5.12	DOD Internet Protocol Hierarchy	5-83
5.13.3	Composite OSI Model	5-95
5.13.4.2-1	IEEE LAN Reference Model	5-99
5.13.4.2-2	IEEE 802 LAN Options Available	5-100
5.13.4.2-3	IEEE 8-2 - Three Access Methods	5-101
5.14.2	FDDI 100 Mb/s LAN Topology	5-110
5.15.2	FILAN	5-112
5.16.5.1-1	Maximum Mean Carried Data Rate Versus Actual Transmission Rate (500 Bit Packet and One Active Station)	5-123
5.16.5.1-2	Maximum Mean Carried Data Rate Versus Actual Transmission Rate (1000 Bit Packet and One Active Station)	5-124
5.16.5.1-3	Maximum Mean Carried Data Rate Versus Actual Transmission Rate (2000 Bit Packet and One Active Station)	5-125

LIST OF ILLUSTRATIONS (Continued)

<u>Figure</u>	<u>Title</u>	<u>Page</u>
5.16.5.1-4	Maximum Mean Carried Data Rate Versus Actual Transmission Rate (500 Bit Packet and 100 Active Stations)	5-126
5.16.5.1-5	Maximum Mean Carried Data Rate Versus Actual Transmission Rate (1000 Bit Packet and 100 Active Stations)	5-127
5.16.5.1-6	Maximum Mean Carried Data Rate Versus Actual Transmission Rate (2000 Bit Packet and 100 Active Stations)	5-128
5.16.5.1-7	Mean Packet Delay Versus Number of Active Stations	5-129
5.16.5.2-1	Delay-Throughput Characteristics for Token Ring and CSMA-CD Bus	5-131
5.16.5.2-2	Delay-Throughput Characteristics for Token Ring and CSMA-CD Bus	5-131
5.16.5.2-3	Delay-Throughput Characteristics for Token Ring and CSMA-CD Bus	5-131
5.17.2.4-1	Model Partitioning Diagram	5-138
5.17.2.4-2	Control Flow Diagram	5-138
5.17.2.4-3	Implemented Protocols Diagram	5-140
5.17.2.5.1-1	Ethernet All Data Throughput	5-142
5.17.2.5.1-2	Ethernet Client Data Throughput	5-143
5.17.2.5.1-3	Ethernet Collisions	5-144
5.17.2.5.1-4	Ethernet One-Way Delay	5-145
5.17.2.5.2-1	Input Versus Throughput Rate for CPU Factors 1, 0.5 (Single TCP Connection-Single Node)	5-148
5.17.2.5.2-2	Input Versus Throughput Rate for CPU Factors 0.25, 0.125 (Single TCP Connection-Single Node)	5-148
5.17.2.5.2-3	Throughput Rate Versus Max Queue Length (Single TCP Connection - Single Mode)	5-150
5.17.2.5.2-4	Throughput Rate Versus One-Way Delay Time (Single TCP Connection - Single Mode)	5-150
5.17.2.5.2-5	CPU Speed Factor Versus Throughput (Single TCP Connection - Single Node)	5-151
5.17.2.5.2-6	Channel Utilization Versus Throughput (Multiple Connections - Multiple Access)	5-151
5.17.2.5.2-7	Channel Utilization Versus One-Way Delay Time (Multiple Connections - Multiple Nodes, Fixed 11, 440-Bit Message Size)	5-153
5.17.2.5.2-8	Channel Utilization Versus One-Way Delay Time (Multiple Connections - Multiple Modes, Fixed 512-Bit Message Size)	5-153
5.17.2.5.2-9	Channel Utilization Versus One-Way Delay Time (Multiple Connections - Multiple Modes, Fixed 8-Bit Message Size) ..	5-153
5.17.2.5.3	Network Minimum Header Overhead Pie Chart	5-154
5.18.3	Relationships Between a Local Area Network and the Defense Data Network	5-159
8.2.2.3	An Architecture Model for Distributed Processing Systems ...	8-4
8.3.1	Global OSI Reference Model of Computer Based Information Systems Functional Subarchitectures	8-11
8.4.2-1	Single User Cronus DOS Architecture	8-17
8.4.2-2	Multifuser Distributed Cronus System Environment	8-17

LIST OF TABLES

<u>Table</u>	<u>Title</u>	<u>Page</u>
4.8	Objectives for Simulation/Modeling	4-8
5.1.4.1-1	Assumptions and Requirements Influencing DOD Internet Model	5-12
5.1.4.1-2	DOD Protocol Development Time Table	5-13
5.2.2	Local Area Network (LAN)	5-16
5.2.2.3-1	USAF Near-Term LAN Planning Factors	5-18
5.2.2.3-2	USAF Long-Term LAN Planning Factors	5-19
5.4.3	Modular Operations Center (MOC) Operational Staffing Positions	5-30
5.7.1	Hierarchical Classification of Operating System Services ...	5-47
5.11.1.1	Advantages and Disadvantages of Layering	5-76
5.11.2	Groups Involved in Standards	5-79
5.12	DOD Reference Model Layers	5-84
5.13.1	Layers of Open System Interconnection Reference Model	5-91
5.13.3	Functions of the OSI Layers	5-94
5.17.2.5.3	Protocol Layered Data Rates	5-154
5.18.3	Use of ISO Layers in LAN Design	5-158

SECTION 1.0
INTRODUCTION

Statement of the Problem

Distributed, secure, survivable information systems are needed for the strategic/tactical battlefields to enable Air Force personnel to maintain control over forces, provide intelligence about enemy intentions and capabilities, warn of attacks or hostile actions, help conserve resources and aid with countless other tasks. The key to achieving these capabilities lies in applying distributed systems technologies that are combinable to create an integrated, system-wide command, control, communications and intelligence capability (C³I).

The major technologies needed for distributed C³I comprise distributed processing, distributed data base management, distributed network-wide operating system(s), multilevel security, mixed-media data communications, a suite of networking protocols (applications utilities, host-to-host and subnet), internetworking gateways, wide and local area networks (WAN's and LAN's). LAN's are new high-speed (1-200 Mb/s) bus or ring topology digital communication networks which provide shared communication capacity for a localized area of coverage. Some of these needed technologies have been developed for systems distributed over wide area networks, like the ARPANET and the National Software Works. However, newer local area network technologies are just emerging, have a number of different architectures and protocols and exhibit considerably improved performance/cost characteristics. It is not well understood to what extent solutions developed for wide geographic area networks are suitable for the newer local area networks and the new distributed network-wide operating systems to be used with them.

LAN-based networking protocols will be needed for distributed information processing systems. To date, which ones and how they would differ from the WAN protocols has not been well understood. It is recognized, though, that the applications environment in which these protocols will be employed will be one exhibiting a high degree of heterogeneity. The component elements (i.e., computers, peripherals, terminals), the end systems as a whole and the LAN's and WAN's will probably differ in their architectures, interfaces and protocols.

New technology research and development is being pursued by the Air Force in many facets of distributed processing, such as artificial intelligence and expert systems, distributed data base management, distributed network-wide operating system, multilevel security, networking protocols, LAN's, WAN's and mixed-media for LAN's. In order to be able to integrate these individual technologies together into a cohesive system, the Air Force needs a roadmap to

guide this process. This roadmap should address objectives, services, functions, architecture, protocols and eventual military standards spanning the distributed C³I battlefield systems needs. The study and investigation reported on herein for LAN interoperability has examined many of the issues raised above and presents its results, conclusions, recommendations and identifies areas needing further study and development.

SECTION 2.0

BACKGROUND

Background of the Problem

Improving price-performance availability of microprocessors, semi and custom high density VLSI devices, and high-speed local area networks (LAN's) will enable building new information systems which may be functionally distributed in new C^2 applications. These systems will need to be made secure, interoperable and survivable in both strategic and tactical deployments.

Clusters of data processing resources may comprise an organization's subsystem, configured to perform strategic or tactical command and control operations. This is demonstrated by Figure 2.0. These clusters can consist of varying types of computers with different operating system software. Further, new high-speed, low error rate, low delay and high reliability LAN's will make possible the interconnecting together of these data processing resources in single and dispersed cluster (intra-LAN) and multiple localized cluster (inter-LAN) configurations. Through the use of a distributed network-wide operating system and its companion high level application processing protocols, these dispersed data processing subsystems will be able to be integrated into one coherent, responsive and reliable system.

The LAN's which might be utilized to construct a distributed system for C^2 today vary considerably in the service provided, functions performed, their architecture, protocols employed, media, topology and media access control methods utilized and performance. In other words, LAN's are very heterogeneous.

Some LAN's employ twisted-pair, coaxial or fiber-optic media. Either baseband (single) or broadband (multiple) channelization is employed. The media access method can be Carrier Sensed Multiple Access (with or without Collision Detection), Token Bus or Token Ring protocols. LAN's are characterized by their limited diameter of connectivity (maximum about 10 kilometers), by their very low transmission delay (less than 1 microsecond per bit) and high transmission speed (1-200 Mb/s).

LAN's can provide variable high throughput - low delay, low error rate and direct peer-to-peer communications in the local area. For the most part, most communications traffic will originate and terminate within the same LAN between resources distributed on it. However, some traffic will need to leave an originating LAN and terminate on another LAN. Though physical connectivity can be achieved through the use of a LAN, an integrated C^2 system at a functional level requires a suite of protocols, most residing above the LAN's physical media access protocol level.

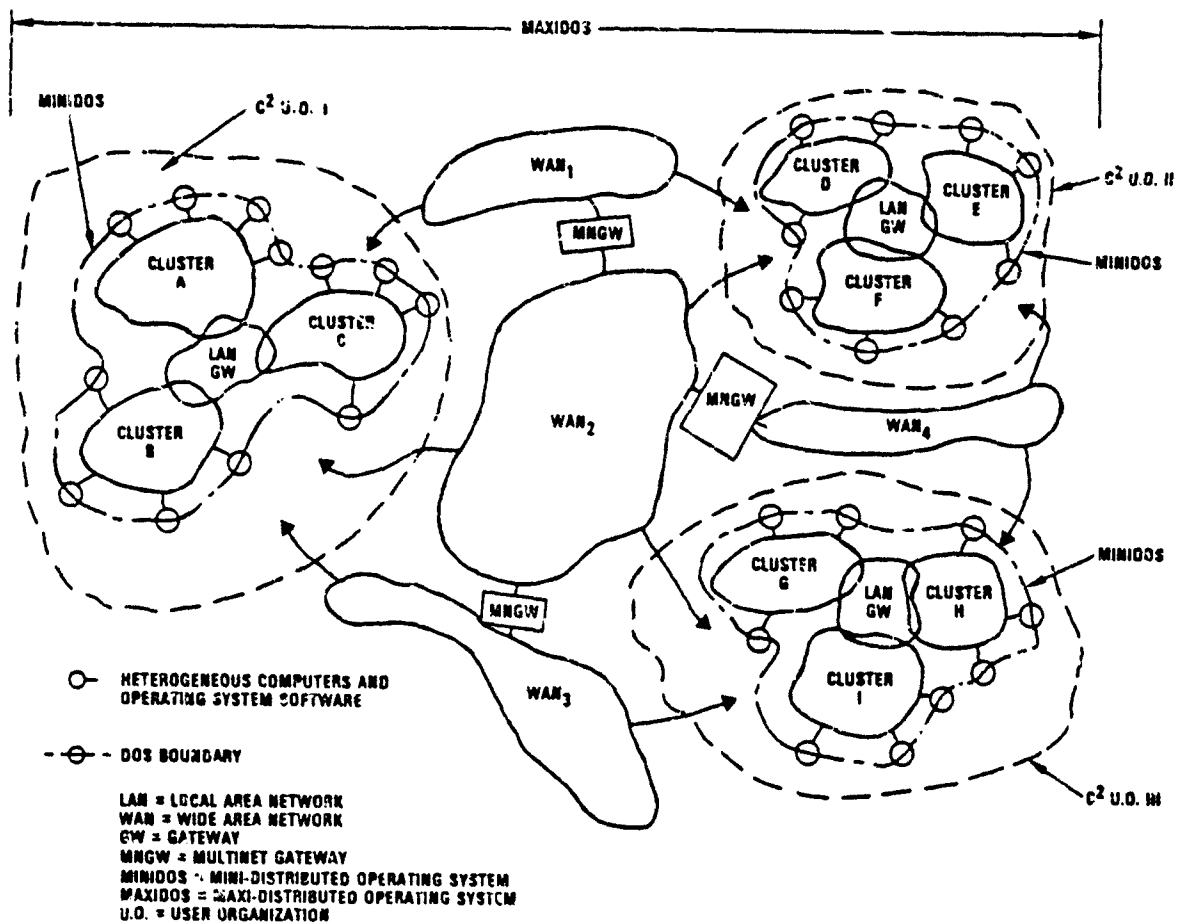


Figure 2.0. LAN-Based Distributed Command and Control Network Model Problem Investigated

In addition to the LAN transmission protocols, the following examples of higher level protocols will also be needed:

- Host-to-host interprocess communications
- File access and transfer
- Terminal handling
- Data base access and update
- Message transfer
- Process-to-process
- Resource monitoring and management

- System fault tolerance/survivability
- Interactive access to distributed processor
- Internetworking (via gateways)

Collectively, these protocols will constitute the procedures or rules by which the various dispersed logical elements of the system coordinate their respective activities to achieve each of the following objectives:

- Interoperability among heterogeneous processing elements within a given LAN
- Interoperability among LAN-based systems, each using different protocols at all levels of abstraction
- Security of information
- Survivability of system functionality
- System-wide control of data processing and communication resources

Geographically dispersed systems, such as the ARPANET [1] and the Defense Data Network [2] have developed and today employ protocols which perform many of the functions discussed above. Those wide area networks (WAN's) and the local area networks (LAN's) of interest for distributed C^2 applications vary widely in their characteristics [3, 4]. It was not clear if these protocols would be adequate (that is, represent an efficient implementation choice) for LAN-based C^2 systems. Further, it was expected that subfunctions within an organization would each be implemented on contiguous but different LAN-based systems.

Protocol layered networking architectures, such as the Open Systems Interconnection Reference Model (OSI/RM) of the International Standards Organization (ISO) [5], provide a formal type of high level abstraction for defining protocols. It was not known whether the OSI/RM's layered protocol structuring contained the appropriate functional definitions within its layers for a LAN-based C^2 system to achieve efficient and effective LAN interoperability.

The following are major issues involved with building distributed information processing systems and do not have satisfactory protocol solutions yet:

- How to make high level applications interoperate compatibly
- What protocols are needed to support high level applications
- Interoperability among LAN-based heterogeneous processing elements
- Interoperability among heterogeneous LAN-based systems internetworked together
- System-wide control of data and communications processing
- Multilevel security and interprocess communications encryption
- Survivability

- Standardized H/W and S/W modular elements
- Subdivision of C² functions into subfunctions distributed over LAN-based processors
- Information sharing
- Directory services of physical/logical objects and resources
- Task scheduling
- LAN-based MININET DOS * (MINIDOS)
- WAN-based MAXINET DOS * (MAXIDOS)
- Cooperating processing
- Device transparency

*DOS denotes Distributed Operating System

The study report's title, "LAN Interoperability," tends to overemphasize the focus which the study conducted. While it was recognized, through close liaison with the study's Contracting Officer Technical Representative (COTR), that LAN's were to be the intended data communications subnetworks for C² systems, the resultant study's focus was steered to be placed upon the information processing user elements which would be attached to a LAN. These LAN users comprise the higher layers of networking protocols and the functions comprising a distributed network-wide operating system. It is within these protocol layers where such networking utility services as file transfer, terminal handling, message exchange and remote job processing are performed by virtual protocols. Exactly how and where these networking utility protocols interface with the distributed network-wide operating system has not been understood and constituted one of the major study issues.

With the DOD having its ARPANET-based networking architecture/protocols and the ISO having its own OSI/RM architecture/protocols, an issue existed at the start of the study relative to whether one, the other or both protocol suites would predominate. More specifically, a question had been raised relative to the continued use by the DOD of its Transmission Control Protocol (called TCP) along with the Internet Protocol (IP). An alternative cited was the National Bureau of Standards Transport Protocol (Class 4), referred to as TP4. Further, while the DOD had issued a directive [6] making use of the TCP/IP protocols mandatory for use with wide area networks, the Air Force, by the start of the study, had gone further by setting a policy [7] requiring use of TCP/IP within local area networks as well.

Several study reports and papers [4, 8, 9] had proposed alternatives to the use of full TCP/IP in a LAN prior to the start of the study. The need existed

to establish an objective quantitative set of data which would set out the performance capabilities achievable using TCP/IP in a LAN, as a basis for technically assessing the impact of complying with the policy directives issued by the Air Force.

The remaining study report discusses the above issues in light of the study's objectives, methodology/simulations employed, results achieved, conclusions drawn, recommendations made and areas identified needing further study.

SECTION 3.0
STUDY OBJECTIVES

3.0 STUDY OBJECTIVES

General

At the outset, the Statement of Work's long-term objectives of the study were to achieve the following within the constraints of the level of effort contracted for:

1. Develop protocols to support efficient and effective communication among application level processes within a Local Area Network (LAN)
2. To investigate the issues associated with interoperability among LAN's based on differing protocols of the transmission through application levels

3.1 Initial Study Investigation Objectives

The proposed investigation objectives were structured to focus on the interprocess communications protocols needed to support higher level application layer protocols. In particular, the primary candidate protocol for consideration was the Transmission Control Protocol/Internet Protocol pair of the DOD, in the context of the IEEE 802 CSMA/CD, Token Bus and Token Ring suite of LAN protocols.

The following initial study objectives were identified to be the major areas to be investigated:

- Selection of baseline command and control system models for the Intra-LAN (clustered MINIDOS) and Inter-Lan (cluster-to-cluster MAXIDOS) configurations
- Establishment of representative data processing equipment/LAN configurations, topologies, functions and work loads within these two configurations
- Review of candidate network architecture, services, functions and protocols needed to support the Distributed Operating System (DOS) high level application protocols and gateways
- Conducting of a series of studies focussing on a localized intra-LAN configuration for the MINIDOS
- Conducting of a series of studies, focusing on a localized set of LAN's interconnected for the MAXIDOS
- Configuring and utilization of a simulation model to enable conducting analysis of four alternative TCP/IP protocol approaches and combinations of commercial/militarized LAN subnets
- Review of the approach taken in the Multinet Gateway design and assessing how the LAN to Multinet interconnect should be done

- Developing of protocols and design approach recommendations reporting on findings from investigating the following specific issues:

Intra-LAN Interoperability (MINIDOS Configuration)

- Suitability of ARPANET protocols
- How LAN characteristics affect application level protocols
- Performance effects of LAN features on application level protocols

Inter-LAN Interoperability (MAXIDOS Configuration)

- Multiple LAN's employing different protocols
- Suitability of WAN gateways to LAN's
- Interprocess communications between different host-to-host protocols

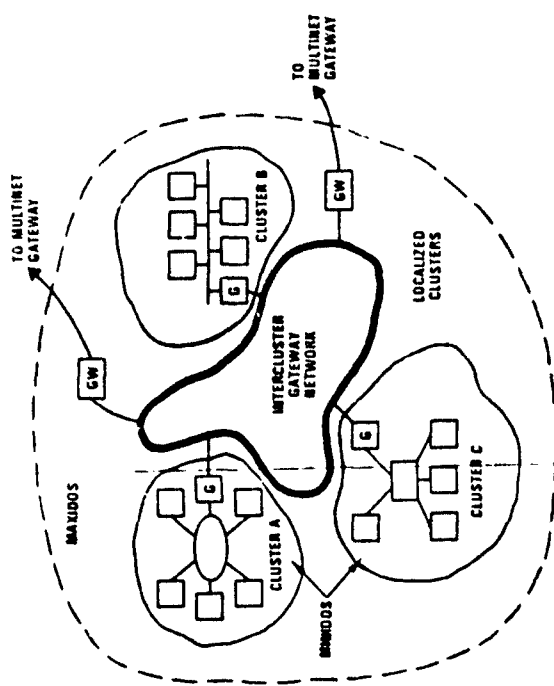
Figures 3.1-1 through 3.1-4 illustrate the MINIDOS and MAXIDOS models of intra-LAN and inter-LAN configurations and the suite of protocols associated with the TCP/IP interprocess communications protocols of interest.

3.2 Revised Study Investigation Objectives

Midway through conduct of the study, the emphasis was shifted away from TCP/IP interprocess communications and detailed performance simulation of multiple LAN protocols to focusing upon the higher layer protocol issues. This occurred after a 6-month review of the progress made to that point.

The revised plan shifted the study resources away from the underlying LAN communications protocols (layers 1-4) and gave primary attention to the information processing region. That is where the higher layers of protocols (layers 5-7) need to become integrated into the distributed operating systems functions to enable building distributed command and control applications processing support services.

To enable a more top down systems driven approach, several tactical and one strategic command and control systems studies previously conducted on requirements and architectures were to be investigated covering the tactical Air Force, Army and strategic DCA missions. Next, conventional and distributed operating systems needed to support the C² mission systems were to be thoroughly studied. The objective here was to examine and establish the structures, major subsystems, services, functions, protocols, mechanisms and issues necessary to support C² systems.



6 = LAN TO LAN GATEWAY
300 = LAN TO WIDE AREA BETWEEN GATEWAY

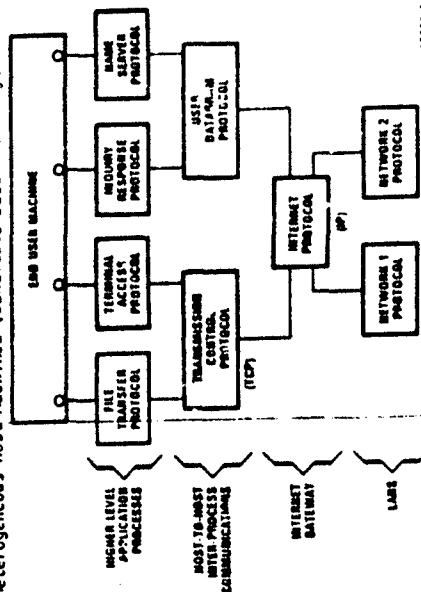


Figure 3.1-4. Protocol Hierarchy Employing TCP/IP

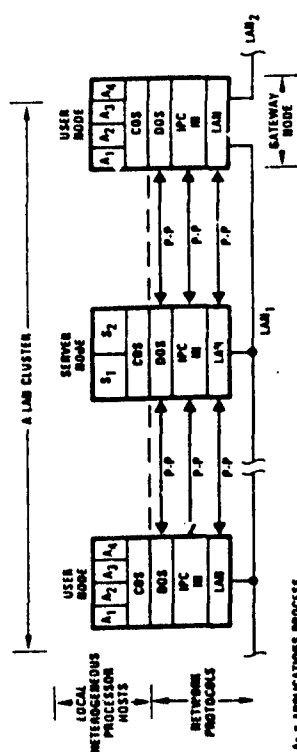


Figure 3.1-1. Intra-LAN (MINIDOS) Cluster of Heterogeneous Host Machines
(Candidate Baseline C₂ System Model)

BASIC LAYER	SUBLEVELS	EXAMPLES
3. APPLICATIONS FUNCTIONS FOR UNLIMITED OPERATIONS SYSTEMS SERVICES	<ul style="list-style-type: none"> • APPLICATIONS PRESENTATIONS • SESSIONS 	<p>FILE TRANSFER PROTOCOL</p> <p>TELECOMMUNICATIONS PROTOCOL</p> <p>INTERNETWORKING PROTOCOL</p> <p>VIRTUAL TERMINAL PROTOCOL</p> <p>REMOTE JOB MANAGEMENT PROTOCOL</p> <p>RESOURCE MANAGER/ PROTOCOL</p> <p>USER DEVELOPMENT</p>
2. END-TO-END NETWORKING AND TRANSMISSIONS SERVICES	<ul style="list-style-type: none"> • HOST-TO-HOST PROTOCOLS (A-B-C, NAME, NETS) • NETWORK TRANSFER PROTOCOLS (A-B-C, NAME, NETS) 	<p>TRANSMISSION CONTROL PROTOCOL (TCP)</p> <p>INTERNETWORKING PROTOCOL (IP)</p> <p>FILE TRANSFER PROTOCOL</p> <p>USER DEVELOPMENT</p>
1. LOCAL NETWORK AND LOCAL NETWORK COMMUNICATIONS SERVICES	<ul style="list-style-type: none"> • PACKET EXCHANGE PROTOCOL • CHANNEL TRANSFER MECHANISM • PHYSICAL INTERFACE 	<p>LOGICAL</p> <p>IEEE 802 LINK CONTROL</p> <p>IEEE 802.3</p> <p>IEEE 802.2</p> <p>IEEE 802.1</p> <p>IEEE 802.5</p> <p>IEEE 802.4</p> <p>IEEE 802.6</p> <p>IEEE 802.7</p> <p>IEEE 802.8</p> <p>IEEE 802.9</p> <p>IEEE 802.10</p> <p>IEEE 802.11</p> <p>IEEE 802.12</p> <p>IEEE 802.13</p> <p>IEEE 802.14</p> <p>IEEE 802.15</p> <p>IEEE 802.16</p> <p>IEEE 802.17</p> <p>IEEE 802.18</p> <p>IEEE 802.19</p> <p>IEEE 802.20</p> <p>IEEE 802.21</p> <p>IEEE 802.22</p> <p>IEEE 802.23</p> <p>IEEE 802.24</p> <p>IEEE 802.25</p> <p>IEEE 802.26</p> <p>IEEE 802.27</p> <p>IEEE 802.28</p> <p>IEEE 802.29</p> <p>IEEE 802.30</p> <p>IEEE 802.31</p> <p>IEEE 802.32</p> <p>IEEE 802.33</p> <p>IEEE 802.34</p> <p>IEEE 802.35</p> <p>IEEE 802.36</p> <p>IEEE 802.37</p> <p>IEEE 802.38</p> <p>IEEE 802.39</p> <p>IEEE 802.40</p> <p>IEEE 802.41</p> <p>IEEE 802.42</p> <p>IEEE 802.43</p> <p>IEEE 802.44</p> <p>IEEE 802.45</p> <p>IEEE 802.46</p> <p>IEEE 802.47</p> <p>IEEE 802.48</p> <p>IEEE 802.49</p> <p>IEEE 802.50</p> <p>IEEE 802.51</p> <p>IEEE 802.52</p> <p>IEEE 802.53</p> <p>IEEE 802.54</p> <p>IEEE 802.55</p> <p>IEEE 802.56</p> <p>IEEE 802.57</p> <p>IEEE 802.58</p> <p>IEEE 802.59</p> <p>IEEE 802.60</p> <p>IEEE 802.61</p> <p>IEEE 802.62</p> <p>IEEE 802.63</p> <p>IEEE 802.64</p> <p>IEEE 802.65</p> <p>IEEE 802.66</p> <p>IEEE 802.67</p> <p>IEEE 802.68</p> <p>IEEE 802.69</p> <p>IEEE 802.70</p> <p>IEEE 802.71</p> <p>IEEE 802.72</p> <p>IEEE 802.73</p> <p>IEEE 802.74</p> <p>IEEE 802.75</p> <p>IEEE 802.76</p> <p>IEEE 802.77</p> <p>IEEE 802.78</p> <p>IEEE 802.79</p> <p>IEEE 802.80</p> <p>IEEE 802.81</p> <p>IEEE 802.82</p> <p>IEEE 802.83</p> <p>IEEE 802.84</p> <p>IEEE 802.85</p> <p>IEEE 802.86</p> <p>IEEE 802.87</p> <p>IEEE 802.88</p> <p>IEEE 802.89</p> <p>IEEE 802.90</p> <p>IEEE 802.91</p> <p>IEEE 802.92</p> <p>IEEE 802.93</p> <p>IEEE 802.94</p> <p>IEEE 802.95</p> <p>IEEE 802.96</p> <p>IEEE 802.97</p> <p>IEEE 802.98</p> <p>IEEE 802.99</p> <p>IEEE 803.00</p> <p>IEEE 803.01</p> <p>IEEE 803.02</p> <p>IEEE 803.03</p> <p>IEEE 803.04</p> <p>IEEE 803.05</p> <p>IEEE 803.06</p> <p>IEEE 803.07</p> <p>IEEE 803.08</p> <p>IEEE 803.09</p> <p>IEEE 803.10</p> <p>IEEE 803.11</p> <p>IEEE 803.12</p> <p>IEEE 803.13</p> <p>IEEE 803.14</p> <p>IEEE 803.15</p> <p>IEEE 803.16</p> <p>IEEE 803.17</p> <p>IEEE 803.18</p> <p>IEEE 803.19</p> <p>IEEE 803.20</p> <p>IEEE 803.21</p> <p>IEEE 803.22</p> <p>IEEE 803.23</p> <p>IEEE 803.24</p> <p>IEEE 803.25</p> <p>IEEE 803.26</p> <p>IEEE 803.27</p> <p>IEEE 803.28</p> <p>IEEE 803.29</p> <p>IEEE 803.30</p> <p>IEEE 803.31</p> <p>IEEE 803.32</p> <p>IEEE 803.33</p> <p>IEEE 803.34</p> <p>IEEE 803.35</p> <p>IEEE 803.36</p> <p>IEEE 803.37</p> <p>IEEE 803.38</p> <p>IEEE 803.39</p> <p>IEEE 803.40</p> <p>IEEE 803.41</p> <p>IEEE 803.42</p> <p>IEEE 803.43</p> <p>IEEE 803.44</p> <p>IEEE 803.45</p> <p>IEEE 803.46</p> <p>IEEE 803.47</p> <p>IEEE 803.48</p> <p>IEEE 803.49</p> <p>IEEE 803.50</p> <p>IEEE 803.51</p> <p>IEEE 803.52</p> <p>IEEE 803.53</p> <p>IEEE 803.54</p> <p>IEEE 803.55</p> <p>IEEE 803.56</p> <p>IEEE 803.57</p> <p>IEEE 803.58</p> <p>IEEE 803.59</p> <p>IEEE 803.60</p> <p>IEEE 803.61</p> <p>IEEE 803.62</p> <p>IEEE 803.63</p> <p>IEEE 803.64</p> <p>IEEE 803.65</p> <p>IEEE 803.66</p> <p>IEEE 803.67</p> <p>IEEE 803.68</p> <p>IEEE 803.69</p> <p>IEEE 803.70</p> <p>IEEE 803.71</p> <p>IEEE 803.72</p> <p>IEEE 803.73</p> <p>IEEE 803.74</p> <p>IEEE 803.75</p> <p>IEEE 803.76</p> <p>IEEE 803.77</p> <p>IEEE 803.78</p> <p>IEEE 803.79</p> <p>IEEE 803.80</p> <p>IEEE 803.81</p> <p>IEEE 803.82</p> <p>IEEE 803.83</p> <p>IEEE 803.84</p> <p>IEEE 803.85</p> <p>IEEE 803.86</p> <p>IEEE 803.87</p> <p>IEEE 803.88</p> <p>IEEE 803.89</p> <p>IEEE 803.90</p> <p>IEEE 803.91</p> <p>IEEE 803.92</p> <p>IEEE 803.93</p> <p>IEEE 803.94</p> <p>IEEE 803.95</p> <p>IEEE 803.96</p> <p>IEEE 803.97</p> <p>IEEE 803.98</p> <p>IEEE 803.99</p> <p>IEEE 804.00</p> <p>IEEE 804.01</p> <p>IEEE 804.02</p> <p>IEEE 804.03</p> <p>IEEE 804.04</p> <p>IEEE 804.05</p> <p>IEEE 804.06</p> <p>IEEE 804.07</p> <p>IEEE 804.08</p> <p>IEEE 804.09</p> <p>IEEE 804.10</p> <p>IEEE 804.11</p> <p>IEEE 804.12</p> <p>IEEE 804.13</p> <p>IEEE 804.14</p> <p>IEEE 804.15</p> <p>IEEE 804.16</p> <p>IEEE 804.17</p> <p>IEEE 804.18</p> <p>IEEE 804.19</p> <p>IEEE 804.20</p> <p>IEEE 804.21</p> <p>IEEE 804.22</p> <p>IEEE 804.23</p> <p>IEEE 804.24</p> <p>IEEE 804.25</p> <p>IEEE 804.26</p> <p>IEEE 804.27</p> <p>IEEE 804.28</p> <p>IEEE 804.29</p> <p>IEEE 804.30</p> <p>IEEE 804.31</p> <p>IEEE 804.32</p> <p>IEEE 804.33</p> <p>IEEE 804.34</p> <p>IEEE 804.35</p> <p>IEEE 804.36</p> <p>IEEE 804.37</p> <p>IEEE 804.38</p> <p>IEEE 804.39</p> <p>IEEE 804.40</p> <p>IEEE 804.41</p> <p>IEEE 804.42</p> <p>IEEE 804.43</p> <p>IEEE 804.44</p> <p>IEEE 804.45</p> <p>IEEE 804.46</p> <p>IEEE 804.47</p> <p>IEEE 804.48</p> <p>IEEE 804.49</p> <p>IEEE 804.50</p> <p>IEEE 804.51</p> <p>IEEE 804.52</p> <p>IEEE 804.53</p> <p>IEEE 804.54</p> <p>IEEE 804.55</p> <p>IEEE 804.56</p> <p>IEEE 804.57</p> <p>IEEE 804.58</p> <p>IEEE 804.59</p> <p>IEEE 804.60</p> <p>IEEE 804.61</p> <p>IEEE 804.62</p> <p>IEEE 804.63</p> <p>IEEE</p>

Figure 3.1-3. Layered Architecture for Communications Network Protocols

Another objective was to concentrate primarily on the LAN-based distributed network operating system and its interface to the localized operating system: the MINILOS and COS.

Additional objectives in the revised study were the following:

- Review of ARPANET higher layer protocols, the DOD TCP/IP protocols and the ISO Open Systems Interconnection protocols, to determine suitability to satisfying the C² and DOS requirements
- Narrowing of the simulation/modeling work to utilize only an Ethernet LAN with TCP/IP and possibly higher layer DOD protocols
- Investigation of internetworking issues associated with joining LAN's to each other directly and by way of a WAN

Since the revised LAN study contract resources would not support protocol development under the revised study scope, the identification of needed C² protocols, services, functions and interface/peer protocol state machines was to be the primary objective. Greater emphasis was to be given to identifying what the issues were, which protocols were needed and what new technology projects needed to be initiated, so as to provide the Air Force a roadmap to guide its technology developments.

SECTION 4.0
STUDY APPROACH

4.0 STUDY APPROACH

Overview

The study dealt with policies, requirements, architectures, protocols and technologies, which spanned all layers of the Department of Defense's and Open Systems Interconnection Reference Models. Additional systems aspects went beyond those, reaching into the heart of command and control systems services, functions and characteristics and those of distributed processing and operating systems.

The early investigations, spanning the first 6 months, primarily focused on structuring the simulation model for later evaluation of the IEEE 802 LAN protocols for Layers 0, 1 and 2, plus the DOD's TCP and IP. At the same time, review of DOD and Air Force policies and architectures for protocols was made. The second half of the study took a more top down systems approach, starting with representative Command and Control requirements leading to examination of distributed network operating systems and the high layer application support virtual service protocols. Simulations of TCP/IP and measurement data later was obtained. Generic gateways and bridges needed for interconnections were studied.

4.1 Assessment of DOD, Air Force and Industry Directions

This part of the study was an ongoing effort to become informed of and to keep apprised of the policies, architectures and protocols likely to become the standards of the overall networking industry groups. At the start of the study, it was apparent that a clear consensus of agreement did not then exist among elements of the DOD, Air Force and industry, regarding a common architecture and the individual protocols to be employed.

4.1.1 Department of Defense Direction

Officials within the DOD's Defense Communication Agency (DCA) were contacted and liaison established during the term of the study. The DCA is the DOD's Executive Agent for the development of protocol standards and has done considerable work on the Transmission Control and Internet Protocols (TCP/IP) [10]. In particular, representatives of the DOD's Protocol Standardization Program [11] and protocol development work were contacted from time to time. This enabled the study to gain access to ongoing work and issues with which those bodies were dealing.

Separate contact was established with the ARPANET's Network Information Center and several sets of the ARPANET Internet suite of protocols were obtained [1].

To assist further in a better understanding of the ARPANET protocols, two members of the study attended a 1-week course on the DOD Internet Protocols. This was given by the George Washington University and was taught by Dr. David Mills, one of the architects of the Internet protocol suite.

About the time of the start of the LAN study, an inquiry had begun over the issue of whether the DOD's TCP protocol should be continued to be used or replaced by one developed by the National Bureau of Standards (NBS). The NBS had developed a Transport Protocol which conformed to the Class 4 protocol of the ISO and was known as TP4. The inquiry was assigned to be conducted under the National Research Council of the National Academies of Science and Engineering. Liaison was established and maintained on an ongoing basis with the administrative office of the National Research Council during the study. It was felt that the outcome of this inquiry might have a profound impact on the future course of the DOD's protocol standardization effort, and on TCP/IP and higher layer protocols in particular.

4.1.2 Air Force Direction

At the start of the study, a member of the Air Force Communications Command (AFCC) informed Harris study members that the Air Force's Air Staff responsible for LAN's had formulated a new policy that "the DOD's TCP/IP protocols are the Air Force standards for connection-based transport and internetwork services within packet-oriented local area networks" [12, 13]. As a result, liaison was established with members of the Air Force's LAN Air Staff office who were consulted periodically to ascertain the Air Force's intent and further policies. A draft LAN architecture document [13] provided more detailed technical insight regarding this policy.

A new Systems Program Office for Air Force LAN's (AFLAN SPO) was established during the conduct of the study. The Director of the office was met and informed about the LAN study's intended objectives. The study team was of the opinion that continuing contact with this office was appropriate since the recommendations the study team might make would possibly have significant impact on LAN protocols and architecture which could affect their use in Command and Control applications.

To assist the study team in keeping apprised of Air Force efforts in developing LAN technology, attendance was made at a 2-day workshop on the Flexible Intraconnect LAN. This was arranged by the Air Force's RADC FILAN Program Office. The current status of the FILAN development, capabilities and future directions were learned.

4.1.3 Industry Directions

The primary sources used for assessing the direction of industry were the international and national standards-making bodies as well as selected vendor product offerings for networking. The ongoing work projects involved in developing the overall architecture and individual protocol standards for the Open Systems Interconnection Reference Model (OSI/RM) [14-18] were the major indicators used to gauge where "industry" was going. The joint agreement by the CCITT and ISO standards-making bodies on a single standard for the OSI/RM occurred during the study period and set a firm direction of the international community on that issue.

During the course of the study, attendance was made at meetings of the IEEE 802 LAN project as well as the ANSI X3T55 High Level Protocols to assess firsthand their work and its possible implications for the issues being studied. Various draft protocol documents and working papers were also obtained as they were considered of relevance to the study. Overall, references [14, 15] were considered the best gauge of industry direction.

Various LAN product offerings by manufacturers were monitored and reviewed to ascertain possible significant impact on future networking directions. In particular, the architecture and product developments reported under way by IBM in its Systems Network Architecture [19], its SNA-to-OSI internetworking [20] and its Token Ring LAN [21] were considered the most significant indicators. In particular, IBM's revealing that its newest additions to its SNA [19] constituted a Distributed Operating System and its embracing of the Open Systems Interconnection architecture and protocols [20] as the way it will enable building global heterogeneous network were considered extremely relevant to the study's main area of concern.

A 1-day seminar on the LAN products offered by Network Systems Corporation was attended. This covered NSC's Hyperchannel, Hyperbus and Netex. A half-day work session was held with the Program Manager of the Multinet Gateway product development with Ford-Aerospace. This provided architectural and interface technical data to the study team of that ongoing RADC-sponsored effort.

The study's Principal Investigator, Mr. Walter L. Elden, was Chairman of the Avionics High-Speed Data Bus Applications and Requirements Task Group (HART) of this SAE-AE9B Subcommittee. A recommendation was made to develop the next generation avionics LAN based on the IEEE 802 Token Ring LAN using

fiber-optic cable, operating in the 50-200 Mb/s. This, along with IBM's announced plan to offer a Token Ring LAN, gave indication that for the longer term, the Token Ring LAN might become dominant.

4.2 Tactical Command and Control System Models and Requirements

Command and Control was considered from both a strategic and tactical frame of reference. Available documentation was more readily obtainable for the tactical deployments to control Air Force air-ground resources and planned Army ground resources. However, the primary references employed were [22, 23, 24]. [22, 23] covered Air Force models, requirements and recommended architectures and [24] focused on an Army system architecture for the electronic battlefield.

In each of these references, the following characterizations were examined and common relationships extracted to provide a profile to drive the succeeding studies from the top down:

- Mission operations
- Mission functions
- Form of management of operations
- Functions to be centralized or distributed
- Modularity of shelterized configurations
- Interconnectivity among modules
- Degree of resource clustering into cells
- Survivability and interoperability aspects
- Application of conventional and distributed processing

In addition to the three cited references, a fourth source [25] was utilized as an indication of 1990's and beyond considerations for inter-service interoperability. This reference was examined for the following characterizations:

- Tactical warfare environment of the 1990's
- Principal requirements for command and control systems interoperability
- Principal architectural considerations
- Proposed architecture for intersystem interoperability

4.3 Distributed Information Processing

This part of the study reviewed material on both requirements for and architectural criteria essential to achieving a Fully Distributed Processing System (FDPS) capability [23, 24 and 26]. Enslow's criteria for an FDPS was used as the basic criteria considered necessary to be present in a system architecture

and found that this was fully consistent with characterizations given in references [22, 23, 24]. In particular, the following characteristics were examined:

- Services to be provided to users
- Subsystems necessary
- Functions to be performed
- Architectural structure and major interface and communications paths
- Aspects of managing system resources
- Heterogeneity of elements comprising the system (local)
- Role of constituent and global operating systems
- Implication of the role of peer protocols in interprocess communications
- Multilevel security issues

4.4 Operating Systems (Local and Global)

This part of the study reviewed the literature on several aspects of operating systems. The two major considerations were for constituent (or local) and global (or distributed) network operating systems. The references studied were [5, 14, 15, 16, 19, 23, 24, 26, 27, 28, 29, 30, 31]. The following characterizations were sought in reviewing this literature:

- Services provided to the user
- Subsystems required
- Functions to be performed
- Architectural structuring of elements
- Object-based system model
- Relationships between local and global system resources
- User interfacing
- Common command language
- Resource management across a distributed environment
- Virtual networking utility protocols for accessing remote resources
- Protocols for interprocess communications
- Use of local and wide area networks
- Use of bridges and gateways
- Multilevel security

Out of these studies came a realization that there were two forms of global operating systems: a network operating system (NOS) and a distributed operating system (DOS). While reference [28] represented the NOS form by the implementation of the National Software Works project, the references for the

Cronus DOS [29, 30, 31], while called a DOS, really exhibited the characteristics of an NOS. This therefore suggested the need to structure a generic form of NOS (called the GNOS in the study) to provide a global view of the type of operating system needed for command and control protocol requirements definition. The GNOS was constructed out of the composite data given in the references cited above.

4.5 High Level Protocols Required for GNOS

A set of virtual (canonical) networking utility type of application layer protocols were identified out of the GNOS study as being required. These were seen as constituting the elements of the GNOS Object (or Resource) Manager entities that were needed for accessing resources or objects remote or not available on a GNOS node's constituent operating system. Examples are jobs, files, terminals, devices, messages, documents and general processing resources.

The main references studied to determine the characterization of the virtual Networking Utility Protocols were [5, 14, 15, 16, 17, 18, 19, 23, 34, 26, 27, 28, 31]. The following were examined to determine requirements for the GNOS Networking Utilities:

- Services to be provided to the GNOS resource managers
- Functions needed to be performed
- Peer protocols needed to provide the utilities
- Support protocols for syntax (presentation service), conversational mode control (session) and interprocess communications (transport connections or connectionless services)

4.6 Networking Architecture and Protocol Suites (ARPANET/DOD and ISO/OSI)

Two major networking architecture and protocol suites were studied: the ARPANET/DOD's [1, 10, 11, 32-40] and the ISO's OSI/RM [5, 14, 18]. Similarities in the two architectures were examined in addition to the services offered to users, functions performed and elements of the peer protocols. In the case of services needed in the local area network partition of the architecture model, the protocol developments under way by the IEEE's Project 802 for LAN's were considered [21, 41]. In addition, liaison was maintained with experts working on these standards bodies to track progress developing protocols. On the issue of whether to maintain the formal layering of protocols as exhibited in the OSI/RM, references [1, 5, 14, 19, 20, 40, 42, 43, 44] were studied and formed a basis for understanding.

4.7 Performance Characteristics of Local Area Networks

Local Area Networks (LAN's), employing various media, topologies, medium access methods and two major logical link user service protocols, exhibit

different characteristics in services offered versus performance, under varying loading conditions. The literature was studied in order to characterize what already had been analyzed and reported by others. The main aspect of LAN's deemed to constitute the mainstream in this regard was the work of the IEEE's Project 802 LAN. This work analyzed the CSMA/CD, Token Bus and Token Ring medium access methods by an impartial working group of experts. References studied were [1, 30, 41, 43, 45-57]. The following characteristics were studied:

- Connection-oriented and connectionless services
- Singlecast, multicast and broadcast features
- Medium access method
- Throughput versus offered load
- Delay versus throughput
- Overhead contributed from protocol layers

4.8 Simulation and Modeling Performance Evaluation of TCP/ICP/LAN Protocols

This portion of the study constructed a discrete-event simulation model of a suite of LAN-based protocols and conducted performance evaluations. The primary reference sources used for the model's architecture was [35, 36, 58, 59]. The model's design was constructed to be very general-purpose with facilities incorporated to enable varying a number of parameters.

The objectives set out for the modeling are given in Table 4.8. In summary, the objective was to model and evaluate LAN and internet configurations of TCP and IP protocols operating with either Ethernet, IEEE 802's CSMA/CD, IEEE 802's Token Bus or IEEE 802's Token Ring, for the LAN cases, and simulated effects for Wide Area Network cases. Later, if resources permitted, an additional set of objectives was to simulate several of the ARPANET/DOD higher layer protocols for evaluation, too.

In particular, throughput and delay performance under a variety of user and configuration conditions were the primary parameters to be measured. A special property, highlighted in reference [58], was the capability to represent different configurations of local resources (i.e, processor speed, etc.) and ascertain the sensitivity effects when variations were made.

Implementation of the simulation design was made on a Harris H-800 machine.

The programs were written in SLAM/FORTRAN and documented in a Progressive Project Document (PPD).

Table 4.8. Objectives for Simulation/Modeling

A. GENERAL

1. Represent operations of DOD/IEEE 802 hierarchy of layered protocols in multiple nodes attached to a single LAN and with an internet of multiple LAN's connected (e.g., employing DDN, X.25 WAN protocol effects).
2. Derive quantitative data on the external (as seen on the media cable) as well as internal (as seen inside a node at special interest points) operational performance of primary functions and events (in terms of timing and resource utilization).
3. Represent selected baseline versions of the protocols and then vary protocols to determine effects within the LAN environment (e.g., subsetting to improve measures of performance or improve a resource utilization).
4. Obtain a quantitative understanding and a data base of what is happening so as to determine cost/benefit trades, which can lead to recommendations to customer.
5. Primary protocols of interest to model are:
 - LAN - Ethernet Blue Book CSMA/CD Layers 1 and 2
 - IEEE 802.3 CSMA/CD Layers 1 and *
 - IEEE 802.4 Token Bus Layers 1 and *
 - IEEE 802.5 Token Ring Layers 1 and *
 - *IEEE 802.2 Logical Link Control Type 1 Datagram, and then Type 2 Connection Service
 - TRANSPORT - Transmission Control and Internet Protocols of DOD TCP/IP
 - HIGHER LAYERS - TELNET, File Transfer and Electronic Mail DOD Protocols as a minimum
 - WAN - DDN, ARPANET, X.25 Effects
6. Variables of Interest:
 - Topology and Media Access Protocols (CSMA/CD, Token Bus, Token Ring)
 - Data Rates (1, 10, 50 Mb/s)
 - Cable Length (0.5, 1 and 2.5 M)
 - Number of Nodes (2, 10, 50, 100, 200, 500)
 - Types of Nodes (Terminal Concentrator, Resource Server, General Purpose Work Station, Computer Host, Gateway)

Table 4.8. Objectives for Simulation/Modeling (Continued)

- Loadings:
 - User Interface Ports (1, 5, 10, 50)
 - User Traffic (quantity, size, frequency, concurrency, TBD)
 - Physical Partitioning into Several Configurations (inboard versus outboard); 1-N processors, interface buses
 - Internal Resources Variable and Measurable:
 - Memory, buffers, processor cycles, etc.
7. Quantitative Objectives From Results
- Media (LAN) Resource Utilization of Capacity (e.g., throughput, delay)
 - Protocol Overhead Effects
 - Internal Processor - Memory Resource Utilization
 - End-to-end Performance Effects at Interfaces to:
 - Layers 0, 2, 4, 7

4.9 Generic Gateways for Interoperability

This part of the study reviewed literature which dealt with architecture and mechanisms which permitted the interconnection of homogeneous and heterogeneous network elements together. The principal references utilized were [1, 3, 4, 5, 10, 14, 20, 21, 36, 41, 59-70].

The study examined the principles of internetworking given in [62] and [60] which applied them to develop a set of generic gateways. These generic gateways span the full set of OSI/RM and DOD/RM architecture models, from the medium to the applications layer. Not only were methods for interconnecting LAN to LAN by use of bridges and gateways in [60, 62] considered, but the larger scope of LAN to WAN and ultimately the interconnection of computer system of vendor A with computer system of vendor B were considered.

4.10 Reporting of Interim Study Results

The results of the ongoing LAN study were reported on briefly each month in the R&D Status Report CDRL. On a larger scale, two key documents were produced. The first [71] presented an interim technical report covering the period of the first 8 months. The second [72] was prepared and presented at the 1984 Distributed Systems Technology Exchange meeting held at RADC. This constituted a preliminary version of the LAN Interoperability Study Report.

SECTION 5.0
STUDY RESULTS

5.0 STUDY RESULTS

This section presents the detailed results of the studies and investigations conducted. The material presented herein is organized for the most part in the same order of presentation as that in Section 4.0, Study Approach.

5.1 Department of Defense Study Findings

5.1.1 Strategic Distributed and Survivable Command and Control

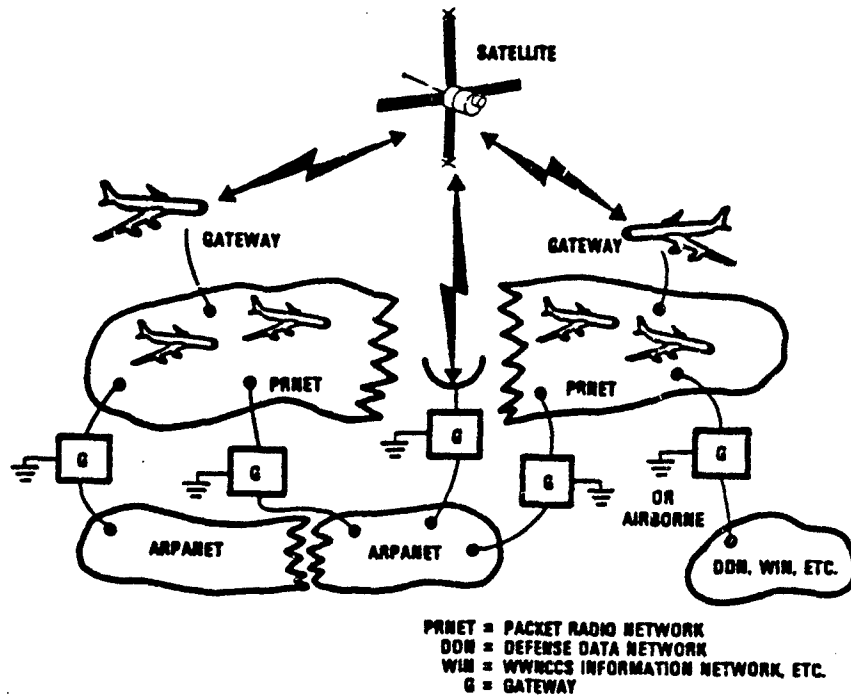
"In military confrontations, the United States forces will face a highly dynamic environment. This environment results from two factors: (1) the increased kill-power and accuracy of modern weapon systems, and (2) the response to this threat - the high mobility of forces during battle. This environment greatly complicates effective command, control, and communications (C³) of U.S. forces. Furthermore, the integration of computer systems and the attendant need for real-time information (data) transfer in a crisis compounds the problems of developing a military structure capable of surviving and functioning in a nuclear engagement" [73].

The Strategic Air Command (SAC) has, as one of its missions, to reconstitute surviving strategic forces for trans- and post-attack command and control. To support this, the DOD is developing new technology for use in survivable C³ for the strategic forces. This is intended for systems that can support an immediate U.S. response to an initial attack as well as meeting a longer-term requirement for C² of surviving forces in response to a protracted nuclear war. Therefore, C³ must survive and offer sufficient flexibility to identify, reconstitute, and employ surviving assets [73, 74].

Through the current deployment of the airborne command post (ABNCP) and the future ground-based, mobile command center, called the Headquarters Emergency Relocation Team (HERT), the Strategic Air Command is attempting to develop survivable C² facilities. As a result, DARPA, SAC, RADC and the Defense Communications Agency (DCA) have agreed and are establishing a test bed to conduct experiments and focus on C³ support to the ABNCP and the HERT. The technologies being developed and evaluated for use in new distributed C³ strategic systems are packet-switching, end-to-end network security and distributed knowledge and data bases.

In a prehostility environment, the present SAC C³ system relies on ground-based strategic data bases. Information is transferred to these data bases by conventional communication systems such as UHF radio, commercial common-carrier systems, and military telephone and land-line data transmission systems. Furthermore, SAC's command and control of its assets are also supported via this

collection of communication hardware. During trans- and post-attack environments, reference [73] states that "under such conditions many of these systems will be at best fragmented; i.e., partitioned in such a way that islands of communications and data base resources exist. In addition, we envision that groups of people will require access to these data bases to carry out their mission effectively." This is illustrated by Figure 5.1.1.



13457-1

Figure 5.1.1. Partitioned Islands of Surviving C³ Resources

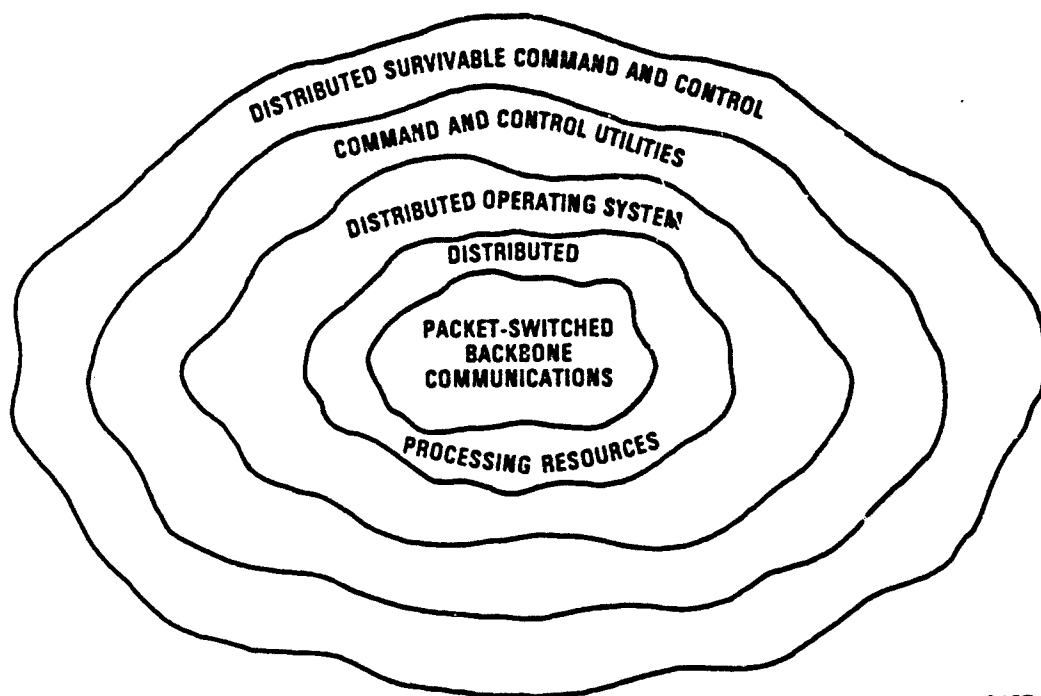
A capability is needed, which does not today exist, that permits the reconstitution of the resources available to these surviving islands, into a unified C³ system. The HERT is the first step toward achieving this goal of providing enduring C².

A systems concept under evaluation for demonstrating such a capability is the following [73]: "First, there is a backbone communication system that can transfer high-speed, nearly error-free data. This system (packet radio) would be deployed on SAC aircraft, to the HERT, and at many strategic locations on the ground. These radios will be networked, providing a means of rapid transfer of information between ground-based and airborne strategic data bases. For survivability, all data bases will be redundantly distributed and will be supported by special-purpose, distributed data base software that ensures that

they contain updated, reliable information. Furthermore, the radios will contain software permitting them to dynamically reestablish communication between airborne users and to reestablish communication with surviving "islands" that would have valuable strategic resources. An automated network management includes the reconstitution of any communication assets that survive the attack (e.g., satellites and ground-based systems that have been augmented to operate in a packet-switched environment" [73]).

5.1.1.1 Command and Control Architecture

"As with the telecommunications and protocol architectures a simple layering of C^2 functionality can be envisioned, as shown in Figure 5.1.1.1. Using the Army as an example, we have many dissimilar processors deployed in a tactical environment. Based on a layered telecommunications architecture, we assume that these resources will be able to communicate reliably (using high level protocols and packet switching) over the backbone telecommunications system. The communications, then, are the nucleus of the C^2 architecture, as depicted in the figure. The next layer of this architecture is made up of the C^2 processing resources that are distributed throughout the battlefield. These processors support not only specific user applications, but also the higher-level telecommunications protocols described above.



13457-2

Figure 5.1.1.1-1. Distributed Processing Layered Architecture for Survivable C^3

In the battlefield, we have a collection of users who need access to these processing resources. Although these users might have primary processors for their use, the C² system architecture should be designed to support an environment whereby backup resources are automatically assigned. Furthermore, user information must be redundantly maintained (for survivability) as, for example, in a "backup cell's" processor. Decisions on where the resources are available to support these functions and on their assignment must necessarily be accomplished automatically if they are to be timely and efficient. To accomplish this management, a distributed internetwork operating system must be developed.

Upon this distributed processing and telecommunications foundation would then be built a collection of generic C² software utilities - which may also be distributed. Examples of such utilities are distributed data base management systems, electronic mail systems, graphics systems and distributed teleconferencing systems.

Finally, at the outermost layer, software systems would be built to support specific C² functions. For example, systems to automate force status reporting and status display are envisioned, as well as systems to support automated sensor correlation and logistics planning [74]."

5.1.2 Interoperability of 1990's Tactical Command and Control Systems

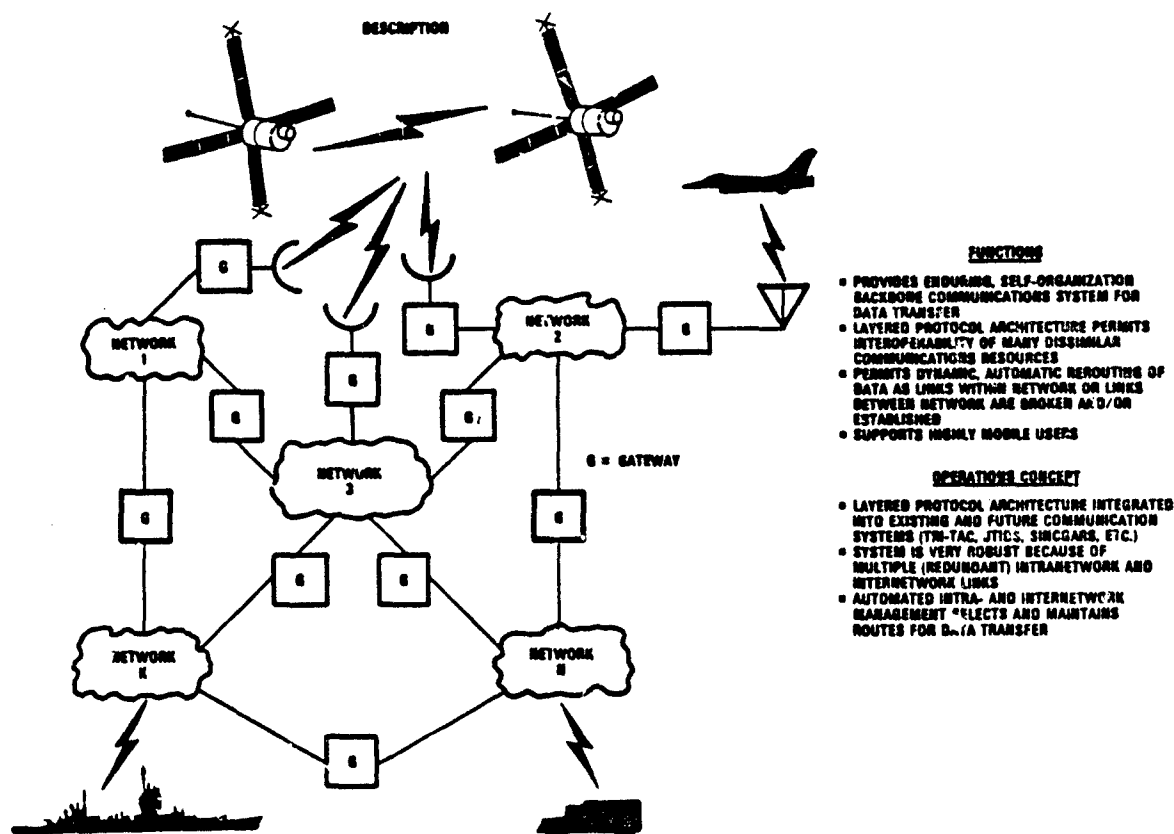
A Joint Service/Office of the Secretary of Defense/Industry Working Group examined, in 1981, the issues involved with tactical information exchange [75] for the 1990's and concluded that tactical command and control systems will need to be interoperable across the military services and their many operational and support systems. This was conducted under the auspices of the Command, Control and Communications Committee of the National Security Industrial Association and was funded by DARPA in conjunction with the Army, Navy and Air Force.

The study considered intra- and inter-service requirements of the services for the year 2000, plus NATO tactical operations. Figure 5.1.2-1 illustrates one example of an automated battlefield set of systems planned for the 1990's needing to interoperate. The tactical warfare environment for the 1990's was considered to comprise dispersed forces with critical nodes, having to operate in all weather, day and night. Non-line-of-sight target acquisition would be needed and indirect weapons delivery. Precise location and status information would be needed. The enemy would be expected to employ communications jamming, exploitation and physical attack.

Characteristics of the proposed architecture for a tactical information exchange system were as follows:

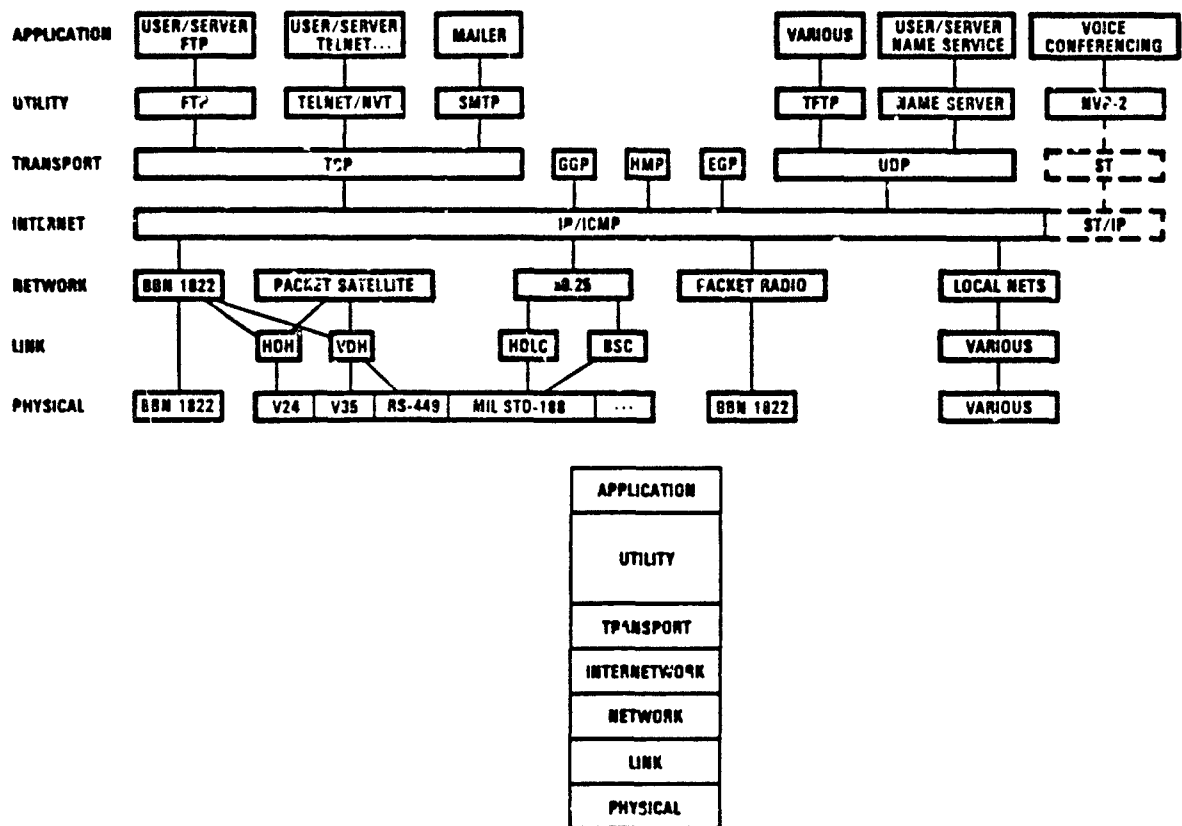
- Packet-switched data backbone
- Standard, multilayered communications protocols for "open systems interconnection"
- Capable of being overlaid on existing link and networks

Figure 5.1.2-2 illustrates the proposed common user's architecture for a tactical information exchange system. The system would interconnect a collection of communication networks for inter-service operability that have implemented the layered protocol model. Figure 5.1.2-3 is an illustration of the DOD's layered protocol architecture model referred to by the study group.



13457-4

Figure 5.1.2-2. Proposed Common User Architecture for Tactical Information Exchange System



DOD INTERNET MODEL

13277-24

Figure 5.1.2-3. DOD Internet Protocol Hierarchy

As a result of the group's study, it made the following recommendations:

1. "Adopt as a long-term policy the evolution of a common-user tactical information exchange structure based on an integrated network of existing, planned and conceptual communications systems. For digital data transfer, the structure should be based on packet-switched networks with specifically defined interconnection protocols.
2. On a high priority near-term basis: complete the definition of a communications protocol reference model that supports DOD needs including those for security, precedence, inter-network data transfer, and support of highly mobile users.
3. When a layered reference model has been established for DOD, protocols should be developed and standardized for each layer of the reference model to support information transfer.

4. Transition plans for steps toward a common-user tactical information exchange system should be developed and implemented both at the individual service level and at the joint service level. The development of a packet-switched data network using multichannel communications equipment for the basic physical link is an example.
5. Special emphasis must be given to the development of network management and control concepts and protocols as they apply to the recommended common-user system, particularly with respect to precedence, security and auditing features.
6. A unified or coordinated test program should be established to support the development and testing of DOD information exchange standards and protocols.
7. Accelerate conceptual and system design efforts relative to the organization, management, control and accessibility of tactical information (at the application and utility/presentation layers of the interconnection model).
8. Current plans for the fielding of tactical communications equipment should be reviewed, with respect to how well they support the recommended future tie architecture framework, this is not a recommendation to stop fielding currently planned communications improvements until the tie framework architecture is implemented.
9. Renewed/increased effort should be addressed to the following:
 - The use of antijammer weapons to off-load communications antijam margin requirements, particularly against airborne jammers
 - The development of deployable communications relay capabilities, particularly airborne
 - The development of multifunction RF systems supporting communications, navigation, identification and surveillance"

5.1.3 Tactical Command Control and Communications for the U.S. Army

"The overall trend toward increasing use of digital data for C² and new Army concepts in survivable command structures implies numerous problems for developing future Army Command, Control, and Communications (C³) systems. For example, the Army's Combined Arms Combat Development Activity (CACDA) developed a concept, called the Cellular Command Post (CCP), that attempts to ensure the survivability of a command center in a tactical nuclear environment through

distribution and replications of the functional areas presently consolidated into one Tactical Operation Center (TOC). In the concept, Division TOC's are divided into 14-16 cells of about 20 persons each. These cells are replicated at least twice, and for survivability they are situated more than 5 km apart.

This TOC architecture raises problems of distributing C^2 information and retaining concurrent replicated data bases at the cells. If assumption of responsibility is to be possible at "backup" cells, then their information must be as current as that used by any (all) other cells. The Division CCP, however, has only a microscopic view of the battlefield. The need for real-time information distribution and maintenance of concurrent data bases in a widely mobile, dynamic, highly dispersed battlefield is a very complex problem.

As a second example, the Army's Training and Readiness Command (TRADOC) has been developing an Army 21 concept. In this concept, the Army of the future is envisioned as: 1) being able to see deep behind enemy lines; 2) being highly mobile and maneuverable; and 3) being able to strike behind the forward line of enemy troops. As part of this concept, it is anticipated that these highly mobile and maneuverable fighting units will, at times, be "communications isolated" from other mobile units or even from higher echelons of command. Under such conditions, commanders will presumably proceed on their own until communications are reestablished.

It is clear from this description that in the Army 21, communications systems that can automatically reconstitute themselves would meet an important need and be of significant advantage. Similarly, processing resources that could function with communications systems that may be intermittently disabled and subsequently reconstituted could provide significant improvements in battlefield information transfer. Specifically, if a mobile, highly maneuverable unit should lose communications, it might proceed to execute a mission and subsequently reappear in an unexpected location. In that case, communications to that unit might not be reestablished efficiently. However, if the communications systems are designed to reconstitute themselves dynamically, then, if any possibility for communications exists, the system will automatically establish the appropriate data paths regardless of user locations. Similarly, if processors are designed to support reliable end-to-end telecommunications protocols regardless of interruptions, then they could, as paths are dynamically established, transfer the data necessary for effective C^2 .

The problem of information management in a tactical environment is even further complicated by the fact that many dissimilar processors are used by battlefield units. Trying to disseminate information reliably between the different machines (which are loosely coupled through various Army telecommunications systems) becomes a classic problem in distributed systems. One solution to this problem would be to converge on a specific machine architecture and system/application software. However, a more pragmatic approach, given the number of computers already in the Army's inventory, is to develop a software and telecommunications environment that supports data transfer between dissimilar processors. In this respect, Army battlefield needs are not so different from commercial user needs" [74].

5.1.4 Department of Defense Protocols Policy

5.1.4.1 Use of Military Versus National/International Standards

Local networks are well suited to intraplatform (vehicle, building,...) applications. Long haul nets (e.g., ARPANET, SATNET, Defense Data Net,...) will be needed for wide-area communications. Packet radio or other mobile digital communication system will be needed in tactical applications involving battlefield automation. No single technology is ideal for all applications, yet the full collection of systems must interoperate.

The military communicator faces a basic dilemma; should military data communications systems use special protocol standards unique to the military, or should they use prevailing national and international standards? References [76, 77, 78] discuss the policy and technical facets of this dilemma.

On this issue, DOD Instruction 4120.20 entitled "Development and Use of Non-Government Specifications and Standards" sets forth the prevailing policy of the DOD concerning adherence to national and international specifications and standards. As applied to protocol standardization, DODI 4120.20 requires standards be adopted as DOD standards in lieu of the development and promulgation of new documents. The instruction does, however, allow exceptions as necessary to provide for unique military requirements. Clarification of the DOD policy by Dr. Richard DeLauer (Undersecretary of Defense for Research and Engineering) was set forth in a memorandum dated 23 March 1982. While reiterating the need to utilize existing national and international standards where possible, he also reaffirmed the current policy of conformance to the existing DOD IP and TCP standards because "military requirements for interoperability, security, reliability and survivability are sufficiently pressing to have justified the

development and adoption of TCP and IP in the absence of satisfactory non-Government standards."

The DOD Internet Architecture Model [77] has evolved over a period of 7 or 8 years, in concert with increasing DOD experience with packet-switched computer communications technology. The model makes full use of the TCP/IP family of DOD protocols. The principal method for achieving interoperability in the DOD Internet Model is the use of a standard Gateway which can route internet traffic from one net to another and the use of a standard set of protocols operating above the internetwork layer. Gateways are specifically intended to support the interconnection of heterogeneous packet nets. The U.S. DOD is moving in the direction of a multi-network or "internet" architecture based on the concept of internet datagrams and gateway interconnection among diverse packet networks.

Table 5.1.4.1-1 lists assumptions and requirements which influenced the DOD Internet Model while Figure 5.1.4.1 illustrates the DOD Internet Protocol hierarchy. This shows protocols above the TCP/IP. DOD has plans for a number of other protocol standards as shown by Table 5.1.4.1-2. These will be done mainly through contractual support and coordination with the NBS. The DOD has developed a working agreement for the development of protocol standards with the NBS through which it will assist the DOD with its technical expertise and also provide representation for the DOD in civilian standards forums.

5.1.4.2 Inquiry Into Use of DOD Versus ISO Protocols

The LAN study learned early of a major issue which developed within the Government over the continued use of TCP and higher layer protocols for networking. The DOD is proceeding with implementation of the Defense Data Network (DDN) using the TCP protocol [2]. At the same time, the National Bureau of Standards has proposed adoption of a Federal Information Processing Standard for the Transport Protocol utilizing the Transport layer specifications developed by the ISO. Both the TCP and the FIPS transport are reported to incorporate the same functional capabilities.

This will become more complex as the ISO is developing a connectionless Internet Protocol which also is reported to incorporate the same functional capabilities as the DOD's IP. Further, both DOD and ISO are developing suites of higher layer protocols for use above the Transport layer.

At the invitation of the DOD, the National Research Council of the National Academy of Science and the National Academy of Engineering set up a committee to examine the layer 4 protocols and make recommendations regarding their use. This study effort began during the summer of 1983 and is jointly sponsored by the DOD and NBS.

Table 5.1.4.1-1. Assumptions and Requirements Influencing DOD Internet Model

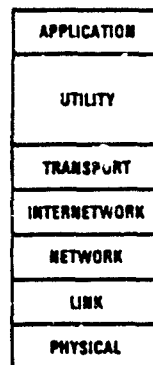
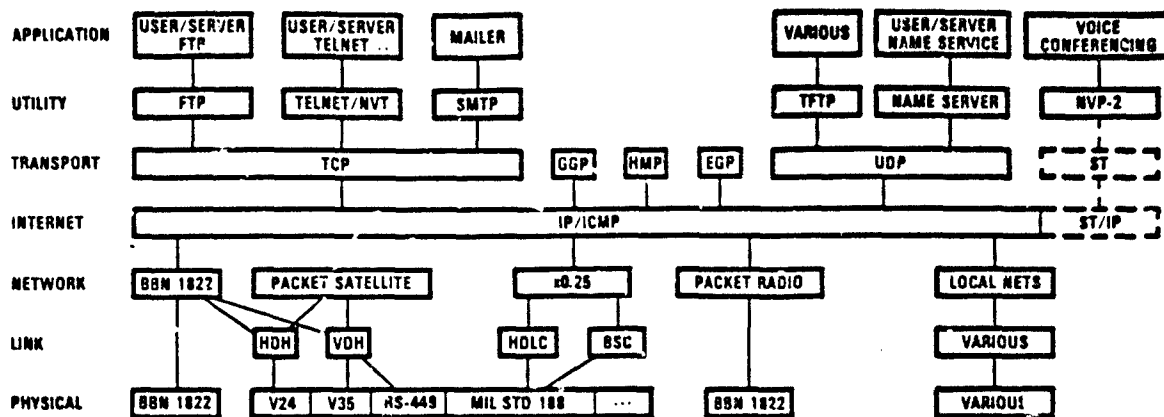
1. Heterogeneous Packet Networks (i.e., Physical, Link, Network Layers differ)
2. Datagram (connectionless) Service at Internet Layer
3. Architectural Provision for Interoperable Tactical and Strategic Communication
4. High Reliability and Survivability Under Hostile Conditions
5. Combined Voice and Data Services
6. Interactive, Real-Time, Transaction, and Bulk Data Transport Services
7. Precedence and Security at Several Layers
8. Broadcast/Multicast Services
9. Host-Host File Transfers and File Access
10. Widely Varying Terminal Types Using Remote Service Hosts
11. Electronic Message Switching Service Utilizing Different Transport Protocols
12. Multimedia (Text, Fax, Graphics, Voice) Electronic Messaging
13. Distributed, Redundant Name-to-Address Translation Services

Mr. Jerome D. Rosenberg, Senior Staff Officer of the National Research Council was contacted several times early during the LAN study to determine progress of the study. The LAN study also learned of, requested, and obtained copies of correspondence [79, 84] between the Computer Business Equipment Manufacturer's Association (CBEMA) and the DOD's Deputy Undersecretary of Defense Research and Advanced Technology over this set of issues.

CBEMA [79] expressed the following request:

"We request that the Department of Defense refrain from issuing requirements for TCP pending a resolution of the protocols question within the Government. We also request that every effort be made to avoid duplication in protocols above layer 4 which have not yet been adopted as military standards."

The DOD's response letter [80] provided the following views (extracted from DOD's letter):



DOD INTERNET MODEL

13277-24

Figure 5.1.4.1. DOD Internet Protocol Hierarchy

Table 5.1.4.1-2. DOD Protocol Development Time Table

Protocol	Initial Spec	Final Spec	Military Standard
Message Protocol (MP)	X	FY 83	FY 84
File Transfer Protocol (FTP)	X	FY 84	FY 84
Terminal Handling Protocol (THP)	X	FY 83	FY 83
Application Level Protocol (ALP)	FY 85	FY 86	FY 86
Presentation Layer Protocol (PLP)	FY 85	FY 86	FY 86
Session Control Protocol (SCP)	FY 82	FY 83	FY 84
Transmission Control Protocol (TCP)	X	X	FY 83
Internet Protocol (IP)	X	X	FY 83
Gateway-to-Gateway Protocol (GGP)	FY 84	FY 84	FY 85
Datagram Network Interface Protocol (DNIP)	FY 83	FY 83	FY 84
X-Task Completed			

"Use of a standard DOD transport layer protocol is sound:

1. DOD military requirement...drives the need for a standard host computer-to-host computer protocol with transport protocol functionality.
2. At the present time, TCP is in the use in many operational environments, but TP (NBS FIPS Transport Protocol) remains a "paper protocol" which has been implemented only in a laboratory environment.
3. TP is not a fully specified protocol.
4. Use of TCP provides DOD system host-to-host interoperability with a proven protocol. There is no alternative commercially supported standard available today.

We have also provided our requirements for protocols above the transport layer to the NBS and are working closely with NBS personnel so that DOD requirements can be reflected in FIPS standards for those higher protocol levels."

The LAN study has maintained liaison with contacts within the Defense Communication Agency (DCA) of the DOD who are providing input for DOD to the National Research Council's study. This liaison will be continued as the Council's inquiry progresses.

At the time of writing this report, the findings and recommendations from the inquiry had not been finalized to allow public dissemination. The results are expected to have a significant impact on future DOD direction, though.

5.2 Air Force Study Findings

5.2.1 Air Force Protocols Policy

The Air Force established in 1983 [12, 13] a policy for packet-oriented local area networks. It states the following as policy:

"The DOD Standard Transmission Control and Internet Protocols (TCP/IP) are the Air Force standards for connection-based transport and internetwork services within packet-oriented local area networks."

This policy and other LAN-related policies as they are developed are to be incorporated in the USAF LAN architecture currently being developed. The Air Force Internetwork architecture is intended to provide a common base within which to provide interoperability among diverse data communities that will exist in the military environment. The datagram-based Internet Protocol (IP) is to allow the flexible evolution of a variety of application level functions, while allowing simple gateways to interconnect long-haul and local area networks. The

Transmission Control Protocol (TPC) is to provide the end-to-end virtual circuit service allowing reliable transmission to take place between subscriber hosts and terminals in a catenet [10].

The following have been identified by the Air Force [12] as pivotal technologies:

1. Inexpensive, powerful microcomputers
2. Inexpensive, high bandwidth communications
3. Proven efficiency of packet switching for bursty computer communications

The key to realizing local area network's potential is with a strategic approach [12]:

1. Build software to span generations of hardware
 - Modular software
 - High level standard languages
 - Hardware insensitive
2. Plan to replace hardware by newer higher performance price offerings
 - Modular hardware
 - Vendor independence at logical interfaces
3. Be willing to pay
 - System performance
 - Software not optimized for speed, memory utilization
 - Initial cost-buy more hardware
4. Mount alternative attack on multilevel security based on encryption, intricately related to protocols
5. Simplify gateways by judicious protocol management

Some choices are considered [12] first order while others are considered second order:

1. Protocols and their management (first order)
 - Layered architecture of vendor independent protocols
 - Choice of protocols, especially at internetwork layer and above
 - Standardization and evolution of these protocols
2. Characteristics of LAN's (second order)
 - Topology (ring, bus, mesh)
 - Medium (twisted pair, coaxial cable, fiber optics)
 - Access mechanism (contention, deterministic)
 - Modulation (baseband, broadband)

5.2.2 Air Force LAN Architecture

The Air Force in 1983 [12, 13] set up working groups to expedite the development on an interim basis the initial development of the USAF LAN architecture and road map. The LAN study contacted the AF Staff office during the summer of 1983, obtained a draft copy of the USAF LAN Architecture and conducted an analysis, presented in the following paragraphs.

Table 5.2.2 presents the current USAF definition for a local area network (LAN).

Table 5.2.2. Local Area Network (LAN)

• Definition*

- A data telecommunications system
- Designed to allow a number of independent devices (e.g., host computers, work stations, terminals, peripherals) to communicate with each other
- Generally, LAN's are restricted to small geographical areas and utilize fairly high data rates
- A LAN is typically a subsystem of a larger information processing system
- Provides functions of data transport, switching and network management
- Excluded functions are higher-layer information processing functions (e.g., file transfer, management information systems)
- LAN's have both physical and logical elements (e.g., cable, media interface, protocol layers)

*Source: "USAF Local Area Network (LAN) Architecture" (Draft), 20 July 1983, HQ USAF/SITT

5.2.2.1 Purpose

The USAF LAN Architecture provides a set of guidelines, policies and standards. It is intended to structure the design, selection, implementation and operation of LAN systems to support user requirements. It addresses such aspects as:

- Security
- Survivability
- Endurance

- Supportability
- Sustainability
- Interconnection
- Interoperation
- Network management/control

It recognizes differences in user requirements and provides for a flexible/modular approach. The ultimate LAN system is considered beyond the current state of the art. The architecture must support technological evolution and be time-phased to enable evolving to the long-term requirements and capabilities.

5.2.2.2 LAN Assumptions

The following assumptions were given:

1. Primary emphasis is to support minimum-essential mission requirements.
2. Support of data and resource sharing at local, regional and global levels required of common-user data telecommunications.
3. Full interconnection and interoperability; both within and between LAN's, are fundamental.
4. Evolutionary approach from near to long term.
5. Network access control, security and auditing capabilities needed.
6. General trend will be toward distributed processing and distributed data bases.
7. Demands will require very high speed (large throughput) data transmission.
8. Value-added network services will be needed (e.g., electronic mail, file servers) to evolve.
9. Evolutionary trend to add video, voice, facsimile to data handling.

5.2.2.3 Near-Term and Long-Term LAN Planning Factors

A set of near-term LAN planning factors was established as well as a set of long-term ones. These are presented in Tables 5.2.2.3-1 and 5.2.2.3-2.

5.2.2.4 Other Aspects of Architecture

At the time the draft USAF Local Area Network (LAN) Architecture document was issued (20 July 1983), several appendix sections were incomplete and were to be provided. They comprised the following:

Table 5.2.2.3-1. USAF Near-Term LAN Planning Factors*

- Commercial technology will determine LAN capabilities initially
- Whether commercial industry incorporates military needs is a marketplace issue
- LAN must support transmission of digital data and support of wideband services preferred
- LAN must support full interconnection and interoperability of an LAN connected device
- LAN's must employ DOD TCP (Full) for transport services and IP for internet services
- LAN's must interconnect with other LAN's and wide area networks (e.g., DDN) using multiple routes and diverse transmission media for survivability endurance
- LAN's must support incremental capacity expansion
- LAN's must avoid single-point failures through distributed designs with functional replication
- LAN's should provide preferential speed-of-service and delivery
- Network access control to evolve to secure LAN implementations
- LAN network interface units (NIU's) should implement protocols through transport layer
- All Internet Protocol (IP) features and services to be supported
- NIU's should be software executable from downline-loadable RAM storage to support network management evolution
- LAN component elements should support functional expansion and growth

*Source: HQ USAF/SITT DRAFT

Table 5.2.2.3-2. USAF Long-Term LAN Planning Factors*

- LAN's must support full range of transmission requirements (e.g., data, voice, facsimile, video, alarms and sensors)
- LAN's must support full interconnection and interoperability between integrated/hybrid information system components
- Dependent upon a universal application of a vendor-independent protocol architecture
- Dynamic networking interconnections and configurations required to withstand stresses
- Full services required to higher-level protocol users of LAN's
- LAN's to support incremental growth in capacity, with fiber-optic systems preferred
- LAN's to support applications (e.g., interactive processing, file transfers, electronic data/voice mail, data/voice/video teleconferencing, workload sharing, distributed processing/data bases, connections, multicart and broadcast)
- Avoid single-point failures
- Provide adaptive resource allocation
- Network access control to support multilevel, controlled and system high modes of operation
- Network interface units perform as full network front ends through virtual terminal service protocols
- Component elements should support expansion

*Source: HQ USAF/SITT DRAFT

1. LAN Protocol Architecture
2. LAN Operation and Maintenance Concepts
3. LAN Network Management and Control
4. LAN Security Architecture

Completed versions of these were not received by the LAN study in time to take them into account.

5.2.3 Air Force Local Area Network Systems Program Office (AFLAN SPO)

A joint Air Force Communications Command (AFCC) and Air Force Systems Command (AFSC) systems program office was set up in 1983 to deal with setting architecture direction for Air Force use of LAN's. At the time of writing this report the AFLAN SPO had not issued its recommendations on LAN architecture or protocols. This will set a very important direction when issued though.

5.3 Industry Study Findings

While the intended use of LAN protocols out of this study is for command and control systems within the military context, the impact of DODI 4120.20, discussed in Paragraph 5.1.4.1 herein, places the consideration of industry protocol standards into relevance. The study, therefore, in addition to examining strategic and tactical requirements and architecture for C² also looked at what industry was doing and where it was going.

5.3.1 Open Systems Interconnection

The primary indicator for industry direction considered was the work under way within the context of the Open Systems Interconnection Reference Model (OSI/RM) being developed within organizations such as ISO, CCITT, ECMA, ANSI, NBS and IEEE. In addition, IBM's System Network Architecture (SNA) was considered.

The scope of the OSI/RM work is worldwide. It is now an international agreement between ISO and CCITT for all future development of standards for worldwide distributed information systems [14]. OSI has been designed to deal with the heterogeneous environment of diverse designs and manufacture. On the other hand, network architecture developed by individual companies, like IBM's Systems Network Architecture (SNA) [20, 82] was designed for a homogeneous environment, where all components are designed and controlled by one organization.

In the command and control environment, exemplified by Paragraphs 5.1.1 and 5.1.2 above, those environments exhibit heterogeneous rather than homogeneous characteristics, and thus, the approach embodied within the OSI/RM seems to be highly relevant. Therefore, DODI 4120.20 and the heterogeneity property of the OSI/RM strongly suggests its consideration to the command and control set of

issues. On the other hand, the OSI/RM has not taken into account the additional special requirements which a survivable and secure [81] military system must deal with. It would be expected, then, that the OSI/RM would lack some functionality needed by C² in those respects. OSI committees have recognized the need for and plan to incorporate security services and functions into the OSI/RM.

The OSI/RM shown in Figure 5.3.1-1, is an abstract description of interprocess communication. OSI is concerned with standards for communication between end systems. In the OSI/RM, communication takes place between application processes running in distinct systems. An end system is considered to be one or more autonomous computers and their associated software, peripherals, and users, that are capable of information processing and/or transfer. Although OSI technologies could be used within a system (and it would be desirable for intra- and inter-system communication to appear as similar as possible to the user), it is not the intent of OSI to standardize the internal operation of each individual system [14].

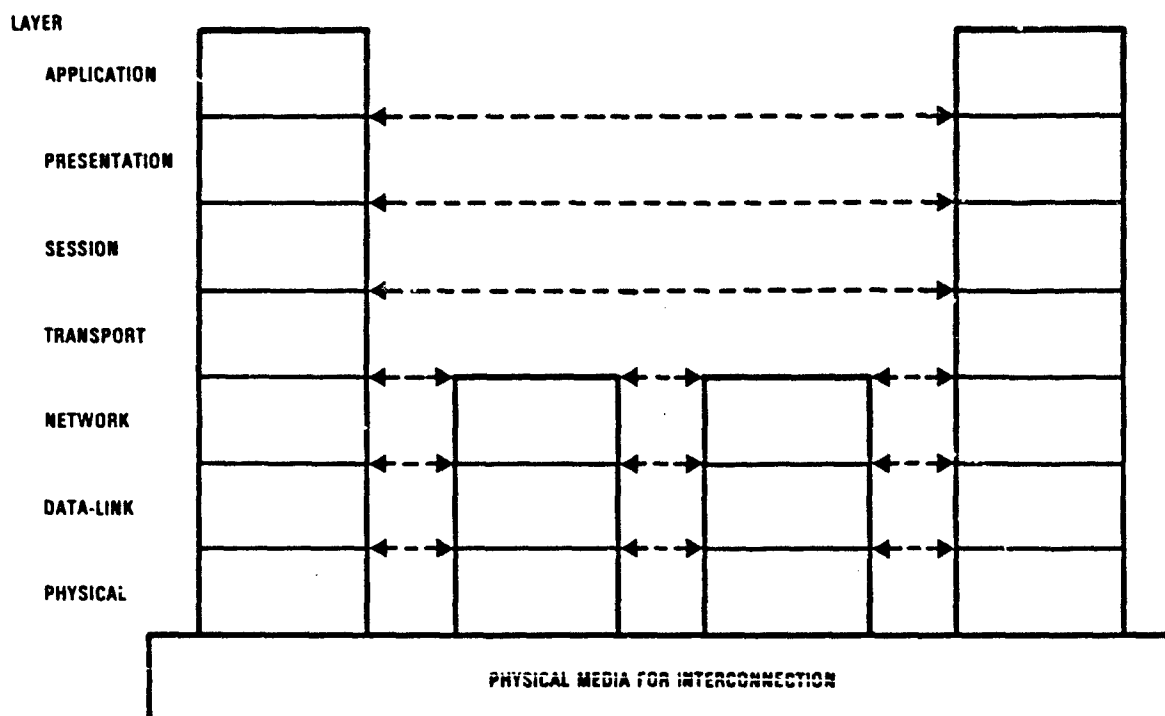


Figure 5.3.1-1. The Seven Layer OSI Architecture

13457-5

OSI is concerned not only with the transfer of information between systems, i.e., transmission, but also with their capability to work together to achieve a common distributed task. In other words, OSI is concerned with cooperation between systems, which is implied by the expression "systems interconnection."

When considering the overall services and functions provided by the seven layers of the OSI/RM, we can split the seven layers into two general functional groups:

- The four lower layers (physical, data link, network and transport), which together support the end-to-end data transport service (i.e., provide the end-to-end communication connectivity). Within this group can reside the use of LAN's, HAN's and Gateways for interworking.
- The three upper layers, which include the application processes themselves, the presentation layer, and the session layer. The latter two provide the necessary means and functions for the communicating entities to organize their cooperation by defining a common syntax for the information exchanged and a means for dialogue scheduling, respectively.

A system which obeys applicable OSI protocol standards in its cooperation with other systems is termed an open system. The objective of OSI is to define a set of standards to enable open systems cooperation among interconnected end systems.

The objectives of a closed architecture, like SNA, and an open architecture, like OSI, can be summarized as follows [20] and illustrated in Figure 5.3.1-2.

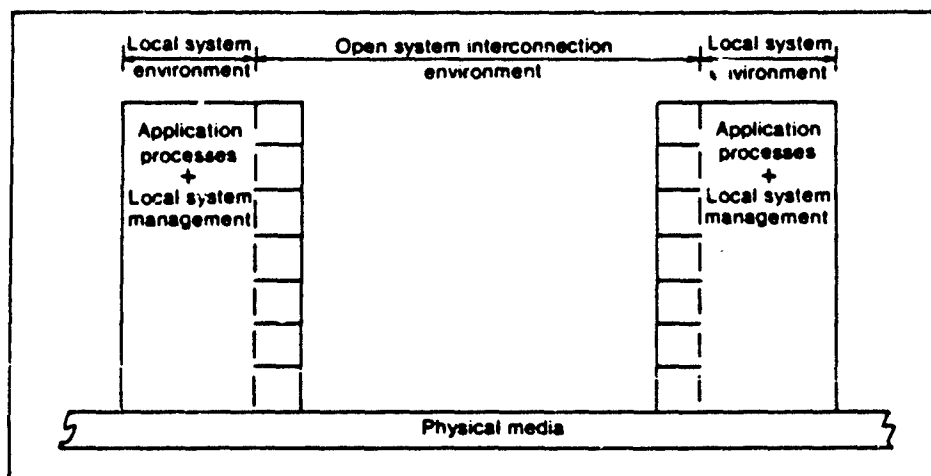


Figure 5.3.1-2. Open Systems Interconnections Between Closed Systems

- "SNA defines the internal structure of a system. It is concerned with the cooperation of products to form a coherent communication system. It also defines the functional responsibilities of each network component within a system, the way information is exchanged between products, and, additionally, the internal control structure that provides the management of system resources and system services.
- OSI defines the external communication capability of a system to make it an open system, i.e., capable of cooperating with other open systems according to OSI standards. The normal OSI applicability is in cases where the cooperating open systems each have a different internal architecture."

"The advent of OSI standards will bring a new dimension to the current environment, by providing universally agreed upon means of permitting communication and cooperation between (or among) heterogeneous systems and products. The existing systems will progressively implement OSI capability in response to user application needs. But the existence of OSI standards should not, and will not, slow down the increasing number and diversity of heterogeneous systems and products. In fact, in response to user's requirements, the systems built on heterogeneous architecture will grow in number and size and they will even provide new functions that are not supported by OSI standards. The resulting OSI environment will therefore be characterized by a large, and possibly increasing, diversity of heterogeneous open systems, heterogeneous because they are built on different architectures; and open because they are capable of cooperating with other systems by implementing the OSI protocols [85]."

One final observation about IBM SNA's future direction and relevance to the LAN Study is given in reference [83]. SNA enhancements now provide Advanced Program-to-Program Communication (APPC). This introduced a new Logical Unit (LU) type, called 6.2., which is referred to as the APPC. For the first time, IBM in SNA states, "that SNA with APPC and other SNA services can be thought of as a distributed operating system." In another reference [84], the following is stated:

"From the perspective of using application programs, SNA is a distributed operating system; it controls the execution of programs and provides those programs with services such as allocation of distributed resources, input/output control, access security, and change commitment control (assurance that all entered commands will be carried out)."

IBM's announced direction of supporting a Token Ring form of LAN for adding to SNA (86) is another important industry indicator.

5.4 Command, Control and Communications for Tactical Air Control System

5.4.1 Tactical Command and Control

Command and control is defined* as:

"The exercise of authority and direction by a properly designated commander over assigned forces in the accomplishment of his mission."

Tactical Command and Control:

1. Command
 - Priorities, strategies, weights of effort
2. Control
 - Match weapons to targets according to command level guidance
3. Communications
 - Internal to the command and control structure including interfaces to strategic communications

Hence, Air Force tactical command control is people, working in accepted and proven military organizations, employing forces in tactical environments - using time proven methods.

5.4.2 Discussion

A large number of tactical automated systems are scheduled for introduction into service in the next decade. These systems will support functions such as sensor data processing, targeting, weapons control and logistics. Emphasis on concepts such as beyond-line-of-sight target acquisition and weapon delivery techniques is generating an increasing requirement for fast, accurate exchange of information across both functional and service boundaries.

Battle environments in which this information transfer must be able to take place will be hostile, dynamic and confused. The system users as well as critical elements of the data distribution system will be dispersed for survivability. Communications connectivities will be continually changing because of jamming, range and intervisibility considerations. Exploitation of signal signatures and content will be a constant threat; in this environment, a shared use, interconnected, multiple link capability is needed to provide efficient, survivable, information transfer services for future C³I users.

*JCS Pub T

ADP users frequently need high data rate communication links, but usually only for short time periods. Because of this bursty nature of the traffic, to furnish each such user with a separate channel would be inefficient allocation of high data rate communications link resources. Use of conventional store and forward data switches would introduce unacceptable time delays. Packet-switched networks, however, support this type of tactical area communications need.

5.4.3 Tactical Air Control Missions and Architecture

A Tactical Air Control system can be defined as the command and control elements that provide the Tactical Air Force Commander with decentralized execution of air space control, area air defense, and air offensive mission control. Mission functions performed would consist of command, weapons control, surveillance, movements and identification, air space/traffic control and battle management. The overall operational management would be centralized, but control of the execution would be decentralized. In order to support such mission operations, a robust/survivable and secure information exchange capability is required.

During the 1980's, many improvements will be made in the Air Force's ability to communicate in a battlefield environment [87]. Programs like JTIDS, SEEK TALK, TRI-TAC, and others will improve the security, jam resistance, connectivity, and capacity of today's Air Force tactical communications. However, some important problem areas will remain, even after these programs have been implemented. Of particular concern are the survivability problems associated with the enemy's modern weaponry.

The goal of the Air Force's communications system planning activities is to reduce or eliminate the remaining problem areas in the 1990's. The basic context and overall C³ architecture for the communications systems planning and system engineering activities are provided by the Master Plan for the Tactical Air Force's Integrated Information System (TAFIIS) [88]. The plan's recommendations call for a dispersed, distributed, survivable C³ system for the 1990's. The plan developed the concept that Tactical Air Control air surveillance and control assets should be configured in a distributed, modular architecture and that the sharing of information was seen as the key to surveillance and intelligence effectiveness. To support the distributed architecture concept of the TAFIIS plan, a study conducted for RADC [89] recommended the need for lightweight, flexible, distributed Modular Operations Centers called MOC's. These MOC's would be interconnected together by way of LAN's to form survivable command operations centers.

The MOC's require the subdivision of the overall TACS airspace control functions (Figure 5.4.3-1) into subfunctions which are then distributed among a configuration of identical command and control shelters (Figure 5.4.3-2), called the Command and Control Modules (CCM's). Interconnectivity between CCM's is provided by means of a standardized, universal bus structure.

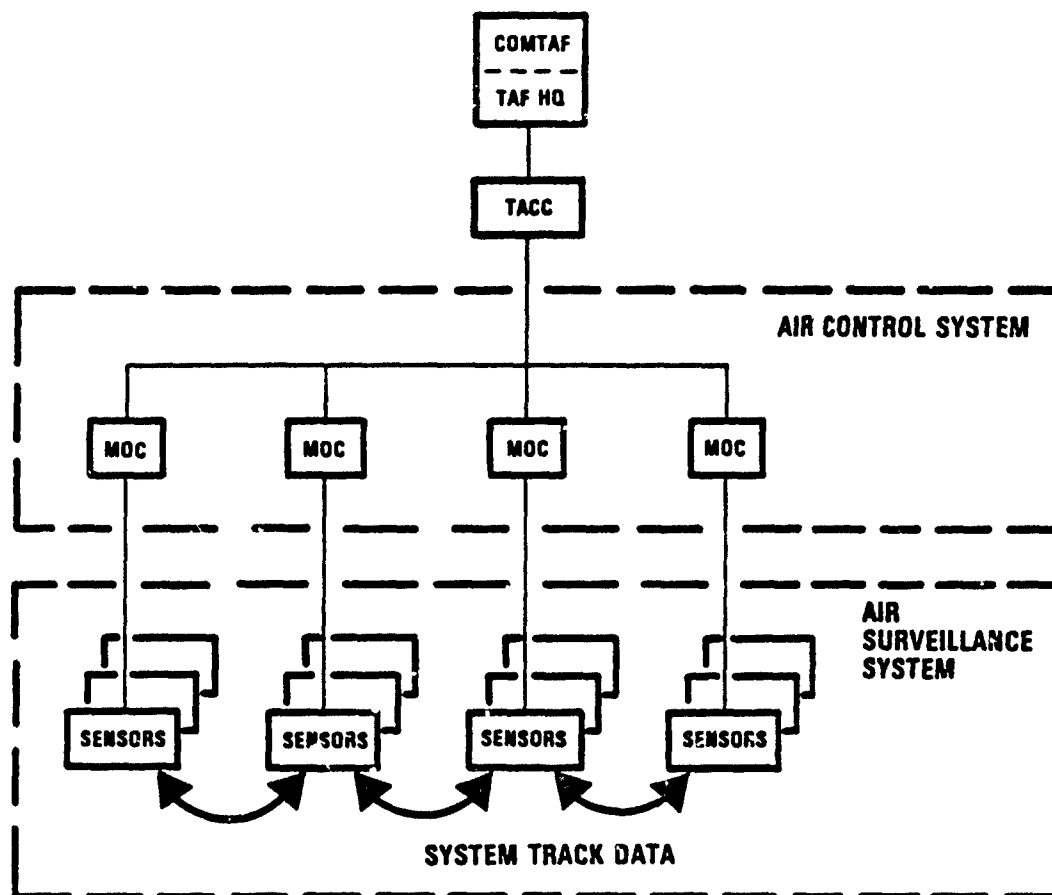


Figure 5.4.3-1. Air Control and Air Surveillance Systems Operational Organization

13277-11

The Modular Operations Centers (MOC's) must be deployable either in a single or multielement configuration (Figure 5.4.3-3). The CCM's must be easily deployable and rapidly relocated to match the rapid pace of projected battles. It was determined that a crew size of four to five people per CCM would be optimum (Figure 5.4.3-4). An S-280 shelter type was selected.

These centers would be interconnected across a battlefield by use of a mixture of tactical Wide Area Network (WAN) facilities. Thus, physical distribution of resources employing networking is intended to reduce vulnerability to detection, increase survivability and an overall flexibility. The study made a

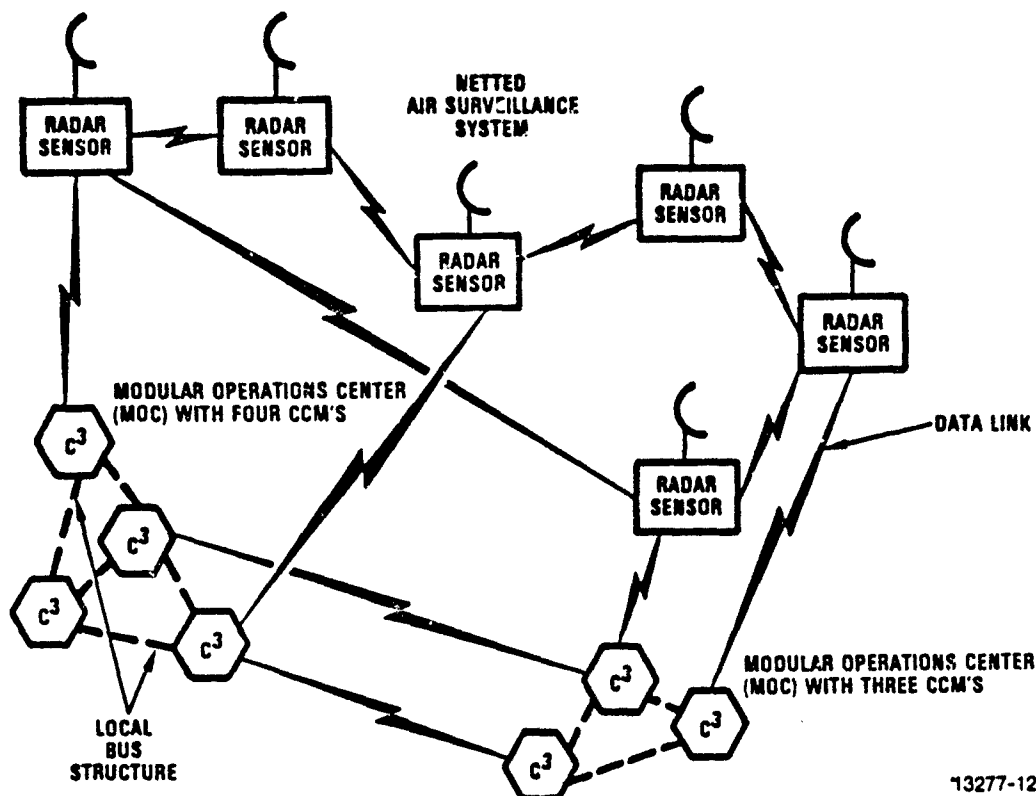


Figure 5.4.3-2. Distributed Air Control and Netted Air Surveillance System Concept

determination of an optimum configuration for these Modular Operations Centers (MOC's), using Command and Control Module (CCM) building blocks. The CCM's are to be designed to employ microcomputer programming to perform the functions of Weapon Control, Movements and Identification, Air Traffic Regulation or Battle Management.

The MOC's should be designed to interface with a netted surveillance system by means of a wideband data bus through microcomputer based interface units in the CCM's. Overall, this is intended to achieve a netted air surveillance system internettted with a netted air control system of Modular Operations Centers.

The CCM would be a general purpose operations shelter that houses processors, integral communications equipment, and operational display positions. One or more CCM's make up a MOC that performs the TAC's air surveillance management and control mission. A CCM is to handle up to five operator positions (Table 5.4.3 and Figure 5.4.3-5). A member of the DOD standard military microprocessor family is to be employed and programmed in the standard DOD higher order language, Ada. Within the shelter, the processing functions are interfaced via an internal bus structure.

MODULAR OPERATIONS CENTER

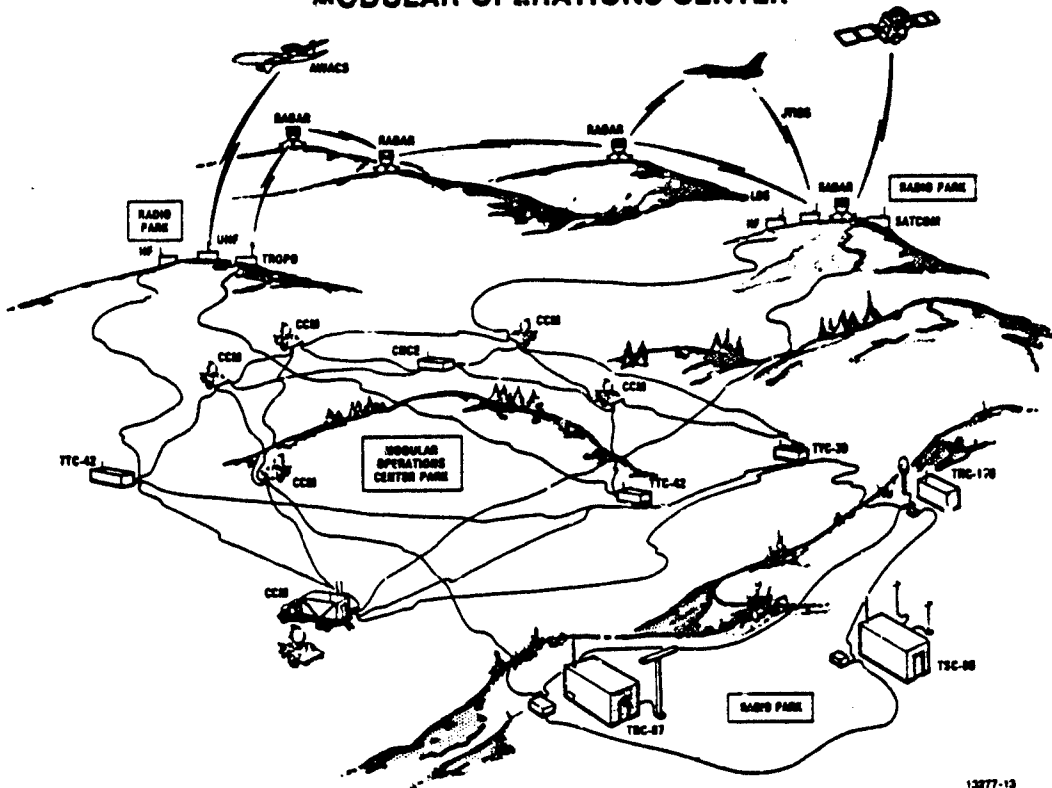
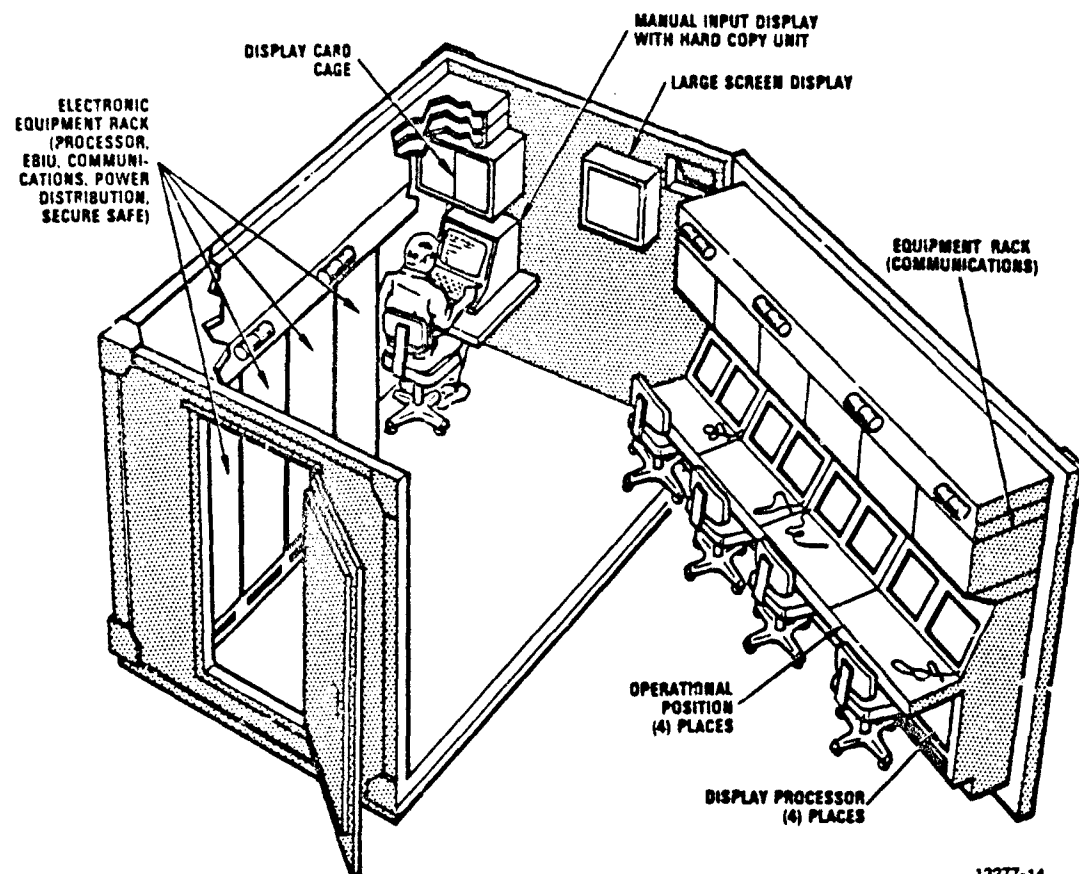


Figure 5.4.3-3. TACS Modular Operations Concept Example Deployment

5.5 Command Control and Communications for the Tactical Army System

A review was conducted of the architectural description [90] of a C³ system to meet requirements of the Tactical Army-Air/Land Battle 2000 concept. The findings were very similar to those recommended in the Air Force TAFIIS Master plan in the following aspects:

- Basic requirement to provide a high degree of survivability and continuity of operations is essential for execution of the air/land battle.
- Intended to be a distributed tactical command and control system.
- Distributed processing was seen needed among the battlefield's automated systems.
- Command posts must be geographically or physically dispersed.
- Certain data bases need to be dispersed and replicated.
- Alternate and backup operations needed.
- System concepts involve the dispersal and interconnection of command cells.



13277-14

Figure 5.4.3-4. Command and Control Module (CCM) Shelter Layout Concept

5.5.1 General Network Architecture

The network architecture is layered to enhance survivability and mobility. The basic subsystems are the subordinate systems connected by the command post network CPN. They constitute the bottom layer.

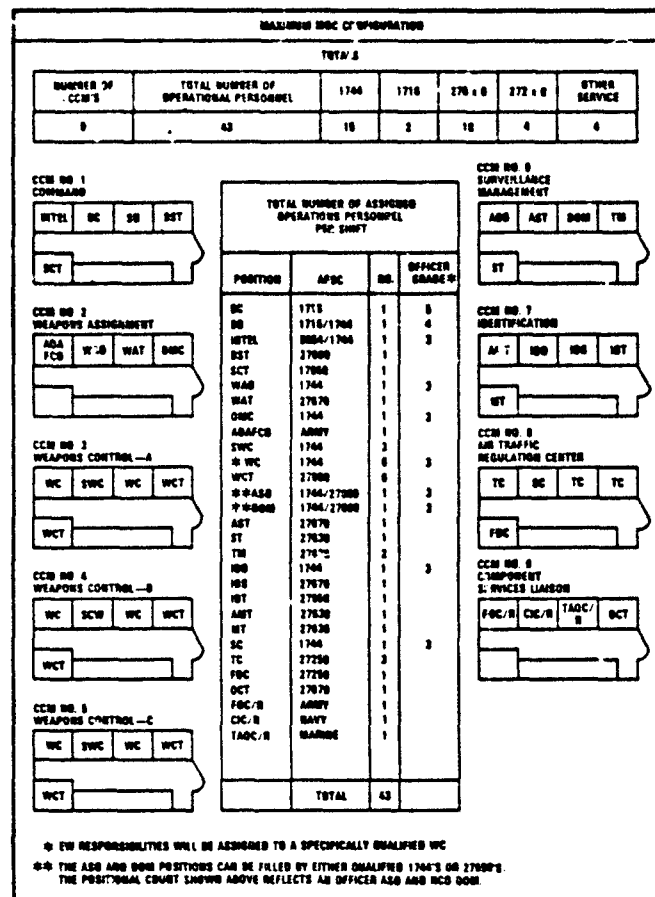
This lowest level subsystem is a cell of the cellular command post and consists of one or more collocated processor elements (e.g., in a single trailer). The CPN is a local area network (max distance between elements 1-10 KM) linking together a set of subordinate systems and a control element. The control element is also a cell of the CCP formed by processor elements. It controls the CPN and provides the gateway capabilities to allow attachment to the backbone network.

The subsystems of the higher layers are configured in the following manner:

Table 5.4.3. Modular Operations Center (MOC) Operational Staffing Positions

<u>Section</u>	<u>Position</u>	<u>Letter Designator</u>	<u>AFSC*</u>
Command	Battle Commander	BC	1716
	Operations Officer	DO	1716
	Senior Director	SD	1716/1744
	Intelligence Officer	INTEL	8054/1744
	NCOIC of Operations	NCOIC/OPS	27690
	Battle Staff Technician	BST	27690
	System Control Technician	SCT	27650
Weapons	Weapons Assignment Officer	WAO	1744
	Offensive Mission Coordinator	OMC	1744
	Air Defense Artillery Fire Coordinating Officer	ADAFCO	Army
	Weapons Assignment Technician	WAT	27670
	Senior Weapons Controller	SWC	1744
	Weapons Controller	WC	1744
	Weapons Control Technician	WCT	27650
Surveillance Management	Air Surveillance Officer	ASO	1744 or 27690
	Air Surveillance Technician	AST	27670
	Data Quality Monitor	DQM	1744 or 27690
	Track Monitor	TM	27630
	Status Technician	ST	27630
Identification	Identification Officer	IDO	1744
	Identification Supervisor	IDS	27670
	Identification Technician	IDT	27650
	Air Movement Technician	AMT	27630
	Manual Input Technician	MT	27630
Air Traffic Regulation Center (ATRC)	Senior Controller	SC	1744/16XX
	Air Traffic Controller	TC	27250
	Flight Data Coordinator	FDC	27250
Component Services Section (CSLS)	Army Flight Operations Center Representative	FOC/R	Army
	Navy Combat Information Center Representative	CIC/R	Army
	Marine Tactical Air Operations Center Representative	TAOC/R	Marine
	Operations Coordination Technician	OCT	27670

*Air Force Specialty Code



13277-15

Figure 5.4.3-5. Maximum Modular Operations Concept (MOC) Configuration

- Subsystems of a given layer are linked together by a network to form a subsystem of the next higher layer. Each subsystem contains a control element that may or may not be distributed. The control element provides management of the subsystem network and an interface to a higher level network.

This layered structure conforms to the Army command structure. The architecture is designed to build subsystems with high intra- but low inter-subsystem traffic.

5.5.2 Subsystem Architecture

The system contains two basic types of networks; the command post network and the backbone network.

- Command Post Network - This is designed around a local area network in a way transparent to its users. It maintains autonomy over its

local functions, can function alone or as an attached element of the overall network. The logical structure is a bus broadcasting network.

- b. Backbone Network - This comprises the three higher layers of the network. Short-, medium- and long-haul networks are utilized.

Each of these networks can function independently and has its own control element. Communications medium will likely be radio, satellite and terrestrial networks, such as TRI-TAC, public network and telephone lines.

5.6 Distributed Information Processing for Command, Control and Communications

5.6.1 General

This subsection reports on two aspects of distributed information processing for C³; what the identified requirements were and what were the recognized characteristics and criteria for structuring a system architecture to satisfy those requirements. Both are discussed in the succeeding paragraphs.

5.6.2 Requirements for C³ Distributed Information Processing

Reference [74], as discussed in Paragraph 5.1.1.1, presented a high level view of an architecture for C³, in that the following broad distributed information processing requirements were stated:

- Command and control functionality should be layered (see Figure 5.1.1.1).
- Many dissimilar processors need to be supported.
- Distributed resources should be capable of communicating over wide area networks employing high level protocols and packet-switching.
- Processing resources will be distributed across the battlefield.
- The processors must support not only specific user applications, but high-level network protocols as well.
- The system must be designed to support an environment whereby backup resources are automatically assigned in the event of a primary's failure.
- User information must be redundantly maintained (for survivability).
- Decisions on where the resources are available to support needed functions, and on that assignment must be accomplished automatically.

- A distributed internetwork operating system is required.
- A set of generic C² software utilization is required, which also may need to be distributed (such as distributed data base management systems, electronic mail systems, graphics systems and distributed teleconferencing systems).
- Automated network management is needed.
- Finally, applications software must be built to support specific C² functions (such as automated force status reporting and support for automated sensor correlation and logistics planning).

5.6.2.1 Strategic Command, Control and Communications Application Characteristics
Reference [34] classified strategic command and control applications into the following:

1. Real-Time Applications

Applications that cannot perform their function if network delay exceeds a certain value. Some applications will fall into this category.

2. Applications That Can Tolerate Delay

This includes interactive applications involving user terminals, where user satisfaction drops when delays increase above very small values. A large number of applications are seen as falling into this category.

3. Applications That Are Relatively Insensitive to Network Delay

This includes batch processing, file transfer and electronic mail applications where throughput is more important than low delay. Network delay typically forms only a minor portion of the total delay experienced by these applications. A sizeable portion of the applications are seen falling into this category.

4. Applications Sensitive to Variations in Delay

Some applications, such as digitized voice, require minimal variation in the delays experienced by the individual packets being received in the data stream. A certain amount of fixed delay is acceptable, but the variation must be held within limits. This will comprise a small portion of the applications. Digitized voice reliability requirements are less than for digital data because information quality is not significantly degraded by small data losses or errors.

Types are classified by [34] as follows:

1. Digital Data

This will be exchanged between the communicating host data processors, the gateways and user terminals/work stations. This data is characterized as follows:

- a. Message Transfer - short transactions of a few thousand bytes or less. Applications include terminal-to-host, terminal-to-terminal, and terminal-to-gateway (distant host) interactions.
- b. Bulk Data Transfer - file transfers, other stream traffic. Applications include data base processing and archiving, and interprocess communications between hosts.

2. Voice - A secure voice capability is seen as being required. This will consist of intercom and teleconferencing. This will require low delay plus controlled delay variance, in-sequencing delivery but not stringent reliability.

3. Video - Video displays and/or teleconferencing will be required. Secure video requires encryption, high data rates and connection-oriented protocol. Absolute reliability for digitized video is not essential.

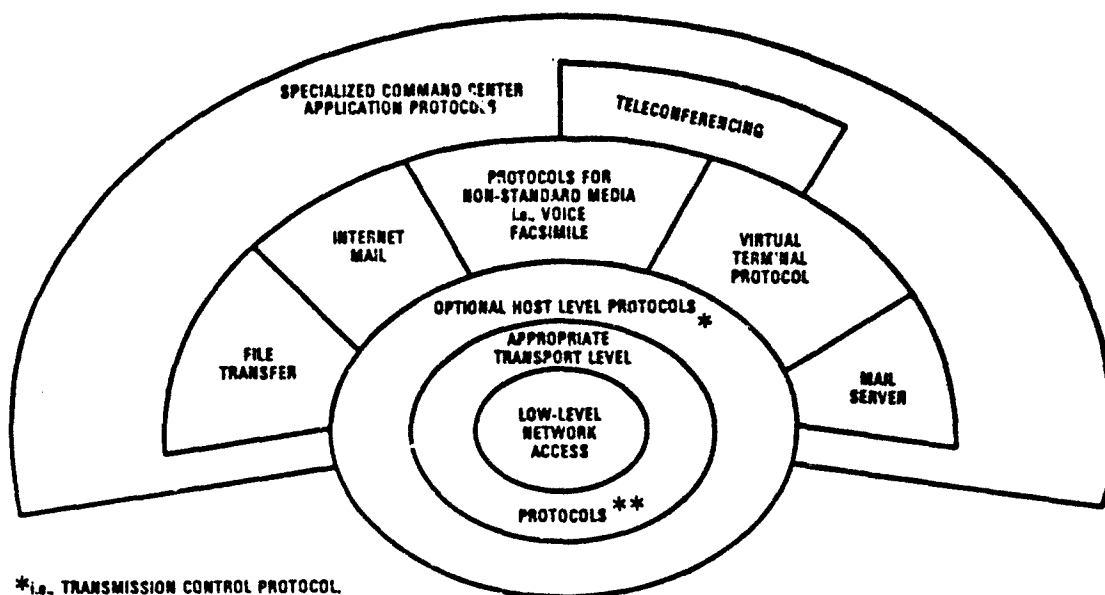
4. Imagery - Included here are facsimile, slow-scan video, and raster graphics. Substantial bandwidth may be required, but the delay requirements are not normally as strict as for video and voice. Reliability requirements are not as strict as those for digital data in terminal or file transfer applications.

5. Non-Digital Devices

A LAN architecture that includes video and voice channels would facilitate the development of multimedia applications. The following are possible ways for incorporating voice and video devices into the LAN:

- Employ separate, dedicated or switched point-to-point channels.
- Digitize analog form information and transmit with the data.
- Employ protocols differing from digitized data to exploit the different reliability and delay requirements.

Several types of networking services and protocols were identified by [34]. The services are as follows, while the main protocols are identified by Figure 5.6.2.2:



* i.e., TRANSMISSION CONTROL PROTOCOL
DISCRETIONARY TRANSMISSION CONTROL PROTOCOL
OR, PROTOCOLS SUITED TO TRANSMISSION OF
NON-STANDARD MEDIA

** i.e., INTERNET PROTOCOL (DATAGRAM)
DISCRETIONARY INTERNET PROTOCOL (DATAGRAM)
BROADCAST

13457-7

Figure 5.6.2.2. Command Center Protocol Architecture

- Terminal-to-terminal communications up to 19.2 kb/s
- Virtualization of terminal characteristics
- Inexpensive terminal interfaces
- Computer-to-computer communication in the megabit range
- High-speed computer interfaces with minimal impact from networking software
- Transaction services (datagrams)
- Virtual circuits (initiation, termination, and control of)
- End-to-end flow control (includes speed matching)
- Error control (ensuring accurate data transmission)
- Equipment interfacing (code translation, terminal translation, etc.)
- Internetworking

This set of services was grouped into three categories, based upon a migration ranging from a basic through a mid-level and reaching an advanced capability. These are as follows:

Basic Requirements

Hundreds of terminals of varying types and tens of heterogeneous computers need to be interconnected. The computers may range from microcomputers to large systems scattered throughout a command center. Local terminal-to-host communications, with speeds up to 19.2 kb/s, will need support. Local

computer-to-computer communications, with data rates of millions of bits/sec, will also need support.

Mid-Level Requirements

High-speed multiplexed interfaces, capable of supporting many varying connections within a host computer, need to be provided. Network-specific software needs to be offloaded from the host computers by use of a host front-end processor. Mechanisms need to be provided to allow access to remote applications, data bases and files. A textual electronic message facility needs to be provided.

Advanced Capability Requirements

Load-leveling and other resource sharing concepts need to be employed to efficiently utilize resources and minimize delay during day-to-day and crisis operations. Name servers need to be utilized to eliminate the need to know the physical location (i.e., net/host) of data bases, standard command center application programs, software tools, or user mailbox. Support is needed for the preparation and transmission of documents with pictorial material, as well as text, and provide high quality printed output. The exchange of critical multimedia message (i.e., voice, pictorial, textual) with multiprecendence levels needs supporting. Lastly, teleconferencing is needed to support interactions among multiple entities.

5.6.2.2 Tactical Distributed Processing Summary Requirements [23]

Automatic data processing requirements include the following minimum areas:

1. Surveillance
2. Command
3. Weapons Control
4. Identification
5. Air Traffic Regulation
6. Other
 - Simulation
 - Support
 - Liaison

Data processing functions fall into the following areas:

1. Composition and editing of data
2. Searching for data records
3. Definition/maintenance of data base
4. Retrieval and output of data
5. User to user data exchange
6. User aids/computation

Components of a TAFIIS Data Processing System Architecture are the following:

1. Processor and peripherals
2. Communication devices
3. User terminals
4. Operating system software
5. System software
6. Application software

Figure 5.6.2.2-1 provides a logical view of the TAFIIS architecture.

Components of a distributed processing architecture are:

1. Users: analyst and system support classes
2. Processing cell of collocated processors
3. MININET network communications interconnecting processors into a cell
4. MAXINET network communications internetting the cells
5. Operating System hierarchical elements:
 - Local (constituent, or COS)
 - MININET functions (or MINIDOS)
 - MAXINET functions (or MAXIDOS)
6. Functional software components
7. Data

Objectives of a distributed processing architecture are:

1. Geographical dispersal of cells
2. Flexible capacity through multiple processors
3. Survivability
4. Security of data and activity
5. Modularity

Sizing is as follows:

1. Cells - tens
2. Users - hundreds
3. Communications
 - Long haul: links hundreds of km long
 - Local: links in the 0.1-1 km

A physical view of the TAFIIS data processing architecture is shown in Figure 5.6.2.2-2.

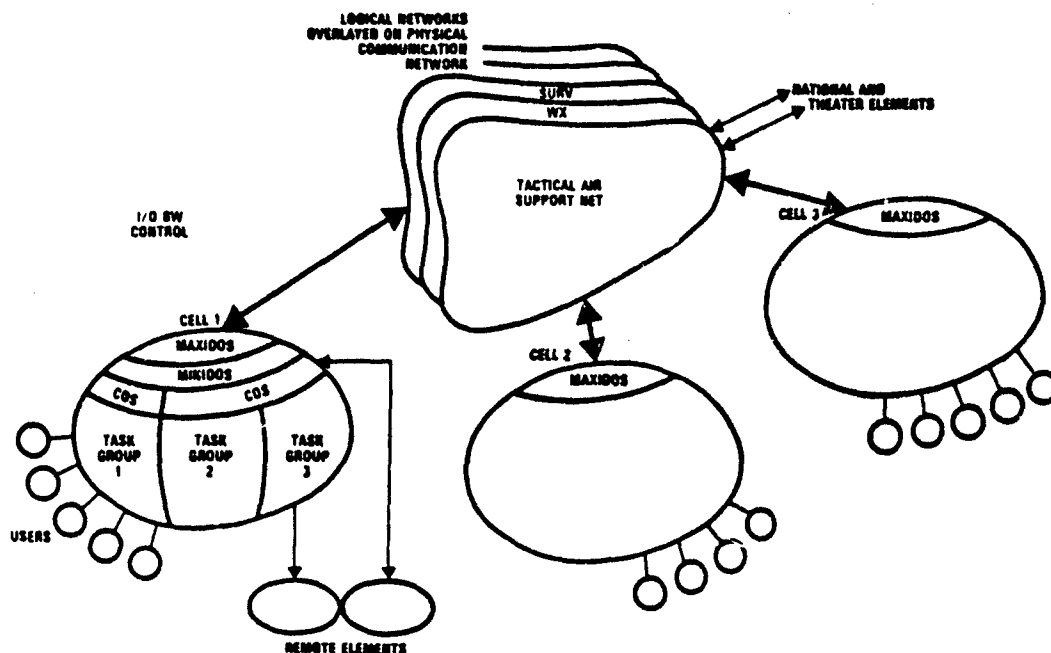


Figure 5.6.2.2-1. Logical View of the TAFIIS Architecture

The overall architecture is as follows:

1. Clusters or cells of tightly coupled processors in a local geographical proximity [23, p. 4-1].
2. Internetting of cells by a relatively low bandwidth network.
3. Geographical distribution of C^3 functions of an automated system.
4. Interconnection of cells is necessary in order to share data and assure survivability of critical information.
5. Choice of tightly coupled local cell of processors based on high common interest among local users in data sharing, processing and information resources in close physical proximity.
6. Choice of loosely coupled internet of cell clusters based on lower level of common need for sharing processing resources and more for accessing a global data base management, categorized by:
 - Data retrieval of information elements between cells
 - Update of other cell elements, particularly for backup purposes
 - Data base segment recovery
7. Cell-to-cell interconnectivity will exhibit outages, errors, and delays to the extent where each cell needs to exhibit a large degree of functional autonomy. This results in the global set of cells forming a "cooperative federation" of processing elements.

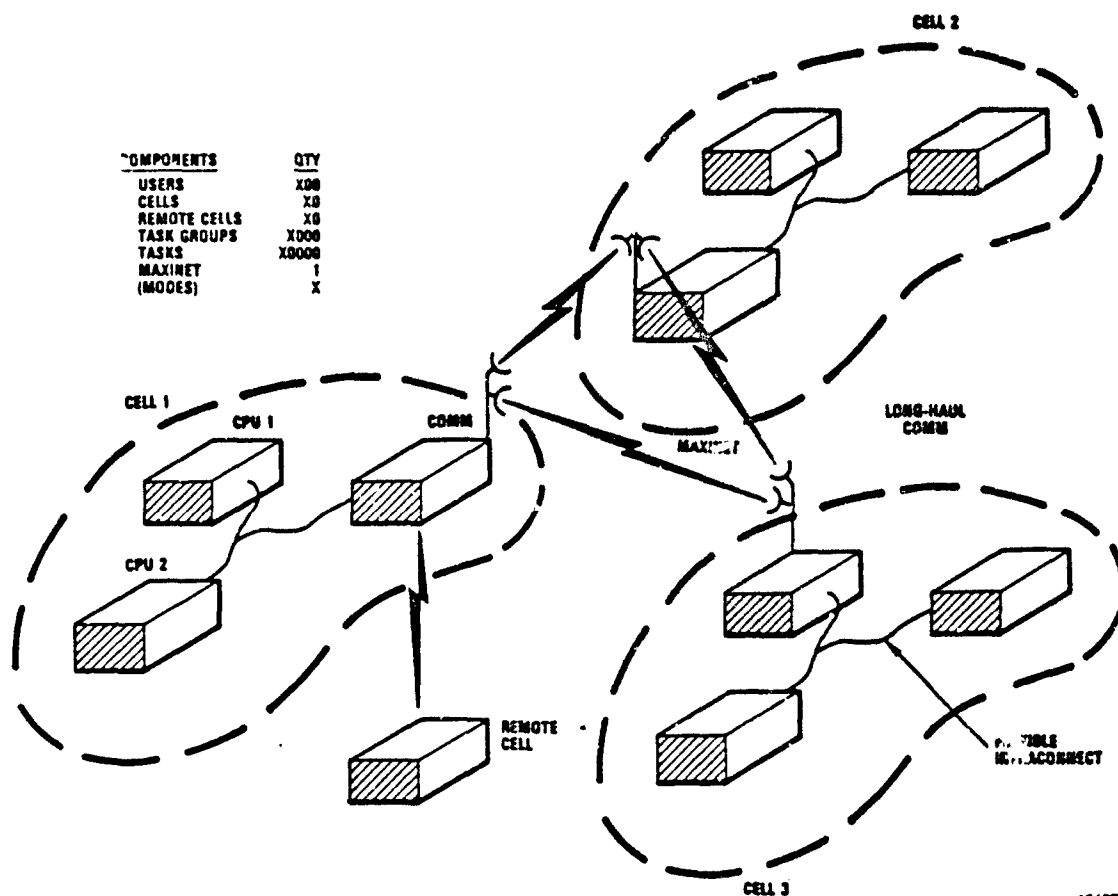


Figure 5.6.2.2-2. Physical View of the TAFIIS Data Processing Architecture

Cell architecture is comprised of the following:

1. A MININET's cell is the unit recognized by the MAXIDOS.
2. Cell comprises all resources interconnected by a single local bus.
3. Hardware components in a cell will be heterogeneous.
4. Program and data within a MININET can be adapted for local performance needs.
5. Programs and data within a MAXINET will have a generic form for transmission.

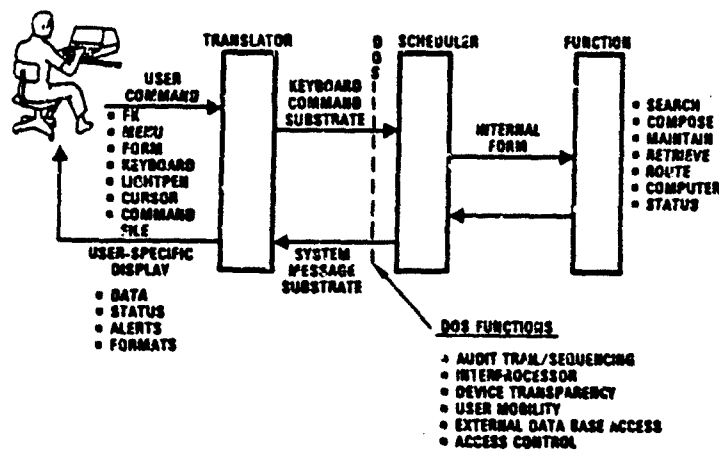
MAXINET logical functions:

1. Addresses the issue of critical information flow and supports system goals of mobility and survivability
2. Recognizes units of MININETS

TAFIIS structured information types would comprise the following:

1. Record messages
2. Data bases
3. User-to-user messages
4. User interface language
5. Software
6. System configuration tables
7. Performance data
8. System messages

The distributed operating system treats all types of data as a message; i.e., a stand-alone unit of information. Figure 5.6.2.2-3 illustrates a user interface architecture.



13467-10

Figure 5.6.2.2-3. Conceptual User Interface Architecture
Resource Management consists of the following:

1. Global resources management via the MAXIDOS through the MAXINET is probabilistic. It operates under conditions of imperfect, incomplete and inconsistent state observations made by a number of loosely coupled MAXIDOS decision makers located in each cell.
2. The TAC C³ environment precludes two or more cells of the MAXIDOS from being able to develop a consensus before an actual resource allocation is made.

3. Each cell's allocation procedures must make independent assessments of the state of the distributed system to arrive at control decisions.
4. Each cell can request status information from other cells to attempt to determine a more accurate estimate of system state.
5. In the absence of good system state status, a cell exercises independent decision making.
6. The resulting control scheme is highly decentralized as there are multiple independent decision makers.
7. Each cell, also, controls only a partial set of the total amount of system (global) resources.
8. At the MININET cell level, the resource management approach is highly dependent on the nature of the mission performed by the cell and has been left unspecified. However, intracell resource management algorithms will be based on control principles requiring high coordination between the allocation procedures resident at multiple processing elements.
9. Therefore, intracell processing will feature centralized control procedures with a resultant ability to assure data consistency and high level adaptive performance of computation tasks.

Security considerations:

1. The tactical system environment requires for the generation, storage, processing and transfer of information elements of diverse levels of security classification.
2. Further, tactical information must be accessed by users at multiple levels of clearance.
3. The proposed security concept relies on these concepts:
 - a. Each cell is responsible for performing user authentication in a well defined, dynamically reactivated procedure.
 - b. Intercell information transfer employs cryptographic protection.
 - c. Cryptographic schemes are varied dynamically with time.

Configuration Management Characteristics:

1. The information processing system needs to deploy a system capable of continuous and efficient reconfiguration of its processing, storage and communications resources.
2. Reconfiguration needs exist both at the cell and global levels.

3. MAXIDOS must be concerned with global management of unreliable resources through unreliable, noisy, low capacity channels.
4. MINIDOS can assume availability of all its local cell resources and can assume a correct view of cell status.

Data Base Management Elements:

1. Data base management responsibilities are broken down hierarchically at the global and local cell levels.
2. At the global level, probabilistic mechanisms are needed.
3. At the local level, more deterministic mechanisms are needed.
4. A high-level information (rather than data) oriented language at the global level is proposed to exchange uniform messages between homogeneous cells.
5. Within cells, translation/conversion schemes are needed to handle heterogeneous processors.

Interprocessor Communication considerations are as follows:

1. Assumed that the intercell MAXINET relies on use of electromagnetic links, which must be shared with other distributed tactical resources. These are subject to failure, damage, interference or destruction.
2. Average MAXINET available bandwidth assumed to be 10-50 kb/s.
3. Cell interfaces with the MAXINET are to present a homogeneous image.
4. Each cell maintains a status map of conditions of other cells and links through the MAXINET.
5. It is desired that intercell communications employ a common language (protocol) designed to support distributed data management functions, such as:
 - Data base element creation/modification/update
 - Data base element availability verification
6. Cells are to communicate a high level by means of a common cell-to-cell language and are not to need to consider the internal operation of other cells.
7. MAXIDOS to adopt an "intelligent gateway" and its common tactical-data oriented language to achieve interoperability.
8. Intracell MININET interprocessor communications need to support interfacing heterogeneous processors. The following will need to be done:

- Reformatting/translation of bit strings elements between heterogeneous processors.
- Reformatting/translation of operating system level requests between heterogeneous processors or, alternatively, translation of such requests to a common interprocessor language.

Protocols [90] are needed to create a virtual network on top of the basic physical network. They will serve two purposes:

- Facilitate the transmission of data and control messages among the subsystem.
- Serve network control in resource allocation, status monitoring configuration and reconfiguration of the network.

Protocols provide for standard information exchange messages between processes.

a. Data Transmission Protocols

These should be compatible with the ISO OSI/RM.

b. Network Control Protocols

1. Configuration Protocols

These support network configuration and two are identified.

- a) Acceptance - used between operating systems when a subsystem is attached to a higher level network.
- b) Exit - used when a subsystem is detached from the higher level network.

2. Status Monitoring Protocols

These monitor the status of the various subsystems and communication links.

3. Contingency Handling Protocols

These are used whenever reconfiguration is needed as a result of orderly detachment of units or as result of a temporary or permanent loss of resources or communication links.

Operational military command and control organizations are looking to local area network technology to provide connectivity among the dispersed data processing elements within their respective organizations. Protocol layers 1-4 provide this basic connectivity. The higher layer protocols, layers 5-7, however, are the ones needed to support the desired functional integration for distributed information processing operations. Protocols are needed for the following data processing functions:

- File transfer
- Distributed data base access and update
- Resource management
- system resource monitoring
- System fault tolerance/survivability
- Interactive access to distributed processes
- Plus others

5.6.3 Distributed Processing Architectural Criteria

Ensrow [26, 91] has set forth some basic criteria which a system must possess in order to be considered a Fully Distributed Processing System (FDPS).

"A FDPS [91] conceptually consists of a loosely coupled network of independent machines. Each machine is capable of communicating with other machines and controls a set of local physical and logical resources (e.g., processors, memory, files, devices, etc.). The machines are autonomous in that each processor or server has final responsibility for the control of the resources it provides. A layer of control is imposed on this network of machines to achieve unification of resources, cooperation, and system transparency. It is assumed that all machines, while retaining their autonomy, follow a common master plan to attain effective cooperation between the loosely-coupled logical as well as physical resources."

The R&D definition has five components:

1. A multiplicity of general-purpose resource components
2. A physical distribution of these physical and logical components
3. A high-level operating system
4. System transparency
5. Cooperative autonomy

Multiplicity of General Purpose Resource Components

System services are provided by a multiplicity of resource components that may be assigned tasks to perform. It is necessary that the system have the ability to be dynamically reconfigured on a short-term basis, with respect to those resources that provide specific services at any given time.

Physical Resource Distribution Through Interconnecting Networking

Physical resources are distributed. They cooperate to provide service to users through the exchange of messages, following rules of protocol. The interconnection comprises networking, connecting peer level resource components

together. The criteria for autonomous operation of processes makes use of information transfers, such as status, requests for service, synchronization between logical resources, and so on.

A High Level Operating System

A well defined set of policies and mechanisms is needed to govern the global operations for the user services and resource providers. There must be no strong hierarchy existing between the high-level and local operating systems. The local operating systems need not be homogeneous.

System Transparency

The interface presented to the user must be one of services and not that of servers. The existence of a distributed operating system is to be totally transparent to the user.

Cooperative Autonomy

The operations of all components or resources, both physical and logical, are to be very autonomous. Message passing network protocols are to be structured to support peer-to-peer interactions, with the right of one to refuse a request for service from another.

In addition, Cypser [92] set forth certain criteria which governs the distribution of functions. First, he identifies six functions which are distributable, to be as follows:

Six Distributable Functions

1. Management of application processing
2. Management of the data that may be stored on a hierarchy of storage devices
3. Management of communications
4. User application programs
5. Application subsystems
6. Input/output mechanisms

Next he states that:

"A distributed function exists if either:

1. The same function can be executed at more than one node in a network and/or
2. The function is not completely executable in a single node and parts of that function are executed cooperatively in separate nodes"

And lastly:

"A distributed function must have the following:

1. Information needed by function
2. Decision logic for function
3. Executability of function
4. Invokability from another node"

5.7 Operating Systems for C³ Distributed Information Processing

One of Enslow's criteria for a FDPS [26, 91] calls for use of a high-level operating system in such a manner as to tie into a global network the collection of individual, probably heterogeneous, local operating system machines. What is an operating system?

5.7.1 An Operating System

The set of programs (software/firmware) that make the physical/logical resources of the system usable is considered the operating system. It is primarily a resource manager for managing processors, storage, input/output devices, and data. An operating system, in addition to managing the system resources, provides services. First, it provides services to internal system users, and, secondly, it provides services to external system users. See Table 5.7.1. Two major forms of operating systems were identified in previous studies as being required to construct a distributed processing system [23, 24, 27, 28, 29, 93]. One is the original nondistributed operating system, called the Local or Constituent Operating System (LOS or COS). The second is the global network-wide operating system. This latter, further, has been viewed as being structured in two different ways; one, call a Network Operating System (NOS) and the other a Distributed Operating System (DOS), and may employ homo or heterogenous OS's.

a. Nondistributed (Local)

This is the existing vendor-supplied OS referred to as a Constituent (or local) Operating System (COS).

b. Network-Wide Operating Systems (Global)

Two forms were identified:

- Network Operating System (NOS)
- Distributed Operating System (DOS)
- They comprise a collection of software and associated protocols that allows a set of antonomous computers; which are interconnected by a computer network, to be used together in a convenient and cost-effective manner.

**Table 5.7.1. Hierarchical Classification of
Operating System Services**

1.0 SERVICES TO PROCESSES

1.1 Process Management

1.1.1 Process Initiation

1.1.2 Process Termination

1.1.3 Error Recovery

1.1.4 Interprocess Mediation

1.1.4.1 Priority Assignment and Management

1.1.4.2 Coordination Primitives

1.1.4.3 Communication

1.1.5 Environment Management

1.1.5.1 Storage Management (for specific processes)

1.1.5.2 Exceptional Condition Management

1.1.5.3 "Privileged" Process Services

1.1.5.4 Process Limit Monitoring

1.2 Resource Management

1.2.1 Processor

**1.2.1.1 Scheduling, Conflict Resolution, Deadlock
Prevention**

1.2.1.2 Allocation

1.2.1.3 Protection

1.2.1.4 Error Detection and Recovery

1.2.2 Timing Services

**1.2.2.1 Scheduling, Conflict Resolution, Deadlock
Prevention**

1.2.2.2 Allocation

1.2.2.3 Protection

1.2.2.4 Error Detector and Recovery

1.2.3 Main Storage Management Global Resource Management

1.2.3.1 Partitioning

1.2.3.2 Segment Control

1.2.4 Secondary Storage Management (Global Resource Management)

**1.2.4.1 Scheduling, Conflict Resolution, Deadlock
Prevention**

**Table 5.7.1. Hierarchical Classification of
Operating System Services (Continued)**

- 1.2.4.2 Allocation
- 1.2.4.3 Protection
- 1.2.4.4 Error Detection and Recovery
- 1.2.4.5 Device Access
- 1.2.4.6 Physical File System
- 1.2.4.7 Process Backing Store
- 1.2.5 I/O Devices
 - 1.2.5.1 Scheduling, Conflict Resolution, Deadlock Prevention
 - 1.2.5.2 Allocation
 - 1.2.5.3 Protection
 - 1.2.5.4 Error Detection and Recovery
- 2.0 SERVICES TO USERS
 - 2.1 System Command Languages
 - 2.1.1 System Operator
 - 2.1.2 Online User
 - 2.1.3 Batch User
 - 2.2 Data Operations
 - 2.2.1 File System Manipulation
 - 2.2.2 Data Generation and Modification
 - 2.2.3 Output Aids
 - 2.3 Program Generation and Invocation
 - 2.3.1 Design Aids
 - 2.3.2 Compilers and Interpreters
 - 2.3.3 Linkers
 - 2.3.4 Library Maintenance
 - 2.3.5 Debugging Tools
 - 2.4 System Management Support
 - 2.4.1 System Generation and Configuration
 - 2.4.2 System Initiation
 - 2.4.3 System Backup and Recovery
 - 2.4.4 Accounting and Auditing
 - 2.4.4.1 Accounting Cost Control
 - 2.4.4.2 Auditing/Surveillance
 - 2.4.5 Performance Monitoring and Tuning

A significant difference in distinguishing between the NOS and DOS forms is based on the degree of autonomous control versus global control. That is, in an NOS all of the processors are independent from a process scheduling point of view, whereas in a DOS the processors do not act independently but in accordance with a global control strategy and some global control synchronization.

Before discussing each of these further, a new concept in structuring of operating systems needs to be introduced. This is referred to as the object model or object oriented design. The following summarizes the object oriented model concept:

Object Based Operating System Architecture

- a. Object oriented design is a design process which deals with abstract (high-level) objects and operations performed on these objects.
- b. Defines an abstract machine composed of abstract resources (objects) that are manipulated by processes to protect the use of resources, to assure coherence, and to impose policies of economic use.
- c. A resource is an abstraction that is defined to the system and given a set of attributes relating to the accessibility of the resource and its physical representation in the system.
- d. A complete operating system is seen as a set of procedure objects and a set of data objects in which the procedure objects represent the actions that can be performed upon the data objects.

Each of these forms of operating systems is discussed next.

5.7.1.1 Constituent Operating Systems (COS)

- a. A local computer system's own operating system; from computer to computer vendor, each would be heterogeneous in characteristics.
- b. COS is the set of services/functions, generally software implemented, which manage the physical and logical system resources.
- c. Serves as a medium level interface between users (people) and the computer's physical machine resources.
- d. Provides services to users:
Process execution control, file manipulation, device manipulation, input/output control, system monitoring/control.

- e. User services provided directly to processes and indirectly via system utility programs.
- f. Resource Managers are programs of the COS which implement the basic policies over a particular class of resources (i.e., CPU, memory, I/O devices, files, communication links).
- g. Resource (device) drivers are physical external resources (e.g., printer). They are managed by special purpose logical functions called drivers.

5.7.1.2 Network Operating System (NOS) [27]

- a. Collection of software services/functions and associated protocols added to heterogeneous constituent operating systems which yield a network-wide set of utility services.
- b. Implemented on collection of large geographic network of computers.
- c. Provides uniform, transparent access to abstract objects and resources distributed around the network.
- d. Objects and resources interact using message passing interprocess communications transactions, loosely coupled.
- e. Location of objects and resources are transparent to users.
- f. Components implemented as ordinary applications software code on Local (Constituent) Operating Systems.
- g. Interprocess communications employ virtual circuits (long message streams) and datagrams (short, interactive transaction messages).
- h. Provides user access to computational services to run programs, manipulate files, input/output data, etc.
- i. Object/resource managers utilize services provided by local operating system, cooperative peer protocols among managers and access networking utilities to use remote system resources.
- j. Objects and object managers:
Files, processes, devices, user ID's, Directory/catalog, etc.

5.7.1.3 Distributed Operating System (DOS) [27]

- a. Collection of new, homogeneous software and protocols which replace previous local operating systems and provide global resource access, allocation and management for users.

- b. Generally employed with localized networks of minicomputers and microcomputers.
- c. Two models generally used for a DOS:
Process model and object model
- d. Process model DOS
 - Each resource is managed by some process.
 - The operating system manages interprocess communications.
- e. Object model DOS
 - Resources consist of objects, which have a type and a set of operations.
 - A capability must be possessed by a user process in order to carry out an operation on an object.
 - Operating system manages the capabilities and to allow operations to be carried out.
 - Interprocess communications mechanism employs either function (or procedure) call or message passing.

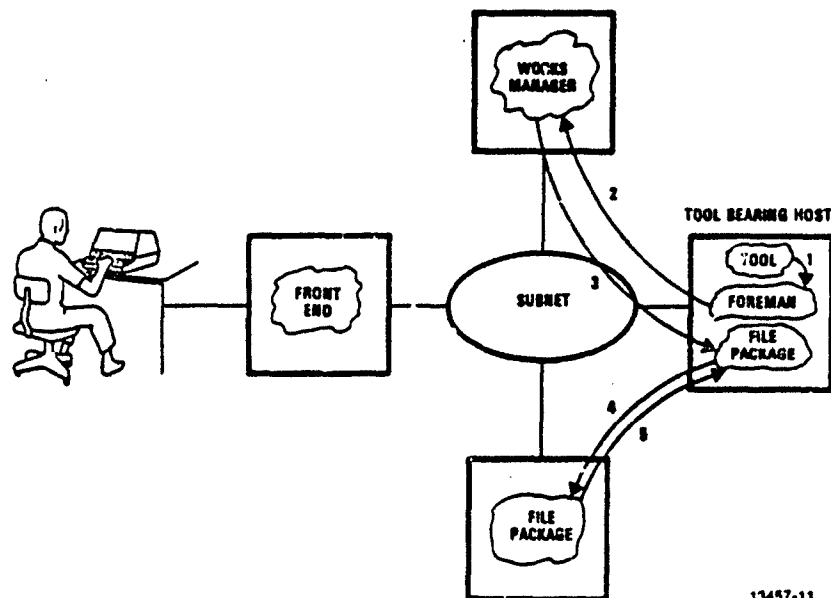
5.8 National Software Works (NSW) Network Operating System

The NSW was reviewed [28]. It represents a good example of a Network Operating System (NOS) as implemented on the ARPANET Wide-Area Network. The NSW provides a network-wide set of utility services to users. The users gain uniform access to NSW managed objects; such as data, files, programs and computing services, that are distributed around the network on heterogeneous machines.

The NSW project applied computer networking technology to improve the distribution of software tools and reusable software modules throughout the DOD. It developed resource management techniques for a large network of mixed-vendor and geographically distributed computers. NSW was one of the earliest and most ambitious attempts at building a NOS.

NSW system components (see Figure 5.8) consist of the following:

- a. Front End process (user interface) provides a command language interpreter
- b. Works Manager (NSW resource manager and access control module)
- c. System Data Base (NSW file system catalog, tool descriptor, user authentication, user rights and privileges)
- d. File Package Processes (responsible for file movement and translation)
- e. Tool Bearing Host Foreman (The tool's interface to NSW)



13457-11

13457-11

Figure 5.8. National Software Worker (NSW) Form of Network Operating Systems

f. Interprocess Communications (message passing mechanism between NSW components)

NSW system implementation is functionally decomposed into units, as follows:

- Works Manager
- Front End
- Foreman
- File Package

NSW components interact cooperatively via message-based interprocess communications along well-defined patterns known as NSW System Protocols. The location of NSW cooperating components relative to each other is completely transparent. Component-to-component interactions are grouped together into patterns called "Scenarios." An NSW scenario consists of the NSW process interactions required to implement a system operation, such as starting a NSW tool or copying a NSW File. NSW components are implemented as ordinary application code on local hosts, whose operating system treats it as any entity making resource demands on the system. NSW system implementation therefore occurs by the programmed cooperation among otherwise independent elements with each local host operating system unaware of the NSW. NSW's Network Operating System consists of

software that was added to the original host computer Constituent Operating Systems. The NSW managed resources are treated as objects. The objects are the elements making up the NSW resource space, such as files and services.

NSW's resource space is defined by the resource catalog. The resource catalog and NSW software using it implement a global symbolic name space for objects. An object's name is like the path through the hierarchical NSW name space. Access control is based on permissions held by the accessing agent referred to as keys. Three kinds of permissions are defined for reading and writing the catalog:

- Lookup (Object lookup)
- Enter (Object name creation)
- Delete (Object name removal)

Other types of permissions are applicable to more object types (e.g., Execute, applicable to service objects. NSW provides public keys; these are available to all users. In addition, there are object attributes:

- Type - Name of system supported types (e.g., File)
- Site - Host on which the objects reside

The following summarizes the main characteristics of the NSW File Systems:

- a. Files are the dominant objects populating the resource catalogs and are key items in the interoperability of NSW program services.
- b. Two forms of storage systems for NSW users and services
 - Long-term, sharable uniform space for files
 - Workspace support is short-term and private
- c. File presentation
 - Data types (text, binary)
 - File structure types (sequential, record)
- d. File Translations
 - Text to text
 - Text to formatted text
 - Sequential text to record structured text
 - Record text to sequential text
 - Record binary to sequential Binary
- e. File Transfers
 - Moved from host to host using connection-oriented interprocess communications

NSW Program Services provides users access to computational services on tool bearing hosts. Services are a type of object managed by the system. Users provided a uniform interface for invoking any service by its resource catalog name (a service spec). NSW services supported are:

- a. Single interactive programs (user is logged into host)
- b. Service bearing host NSW command interpreters (for manipulating NSW files and to run services on host)
- c. Service bearing host native command interpreters (for interacting with standard native command language) to manipulate files and run programs
- d. Service bearing host operating system (initial connection to host)
- e. Batch services (submits, runs NSW jobs, executes and returns results)

Users interface to NSW via The Front End. Front end functions include:

- a. Command interpretation support
- b. Interaction with NSW system components (to run services, manipulate files, etc.) (and status gathering on NSN components and data bases) by way of various protocol scenarios

A communication path employs a Telnet connection between the user's front end and a service. Terminal control is as follows:

- a. User employs a terminal to interact with front end
- b. User employs a Telnet connection to communicate with a local remote front end.

NSW Interprocess Communications was found to have the following characteristics:

- a. Front-End to/from Works Manager (short, infrequent among unrelated processes)
 - User requests for NSW resources and responses
 - Consists of short messages of approximately 1000 bits to and from work manager
 - Transaction processed by any on several works managers
 - Short request, brief delay, short response, a long delay until the next element
- b. Tool/Foreman to/from Works Manager (short, infrequent among unrelated processes)
 - Same characteristics as above, except these are more frequent

- c. Front-End to/from Tool/Foreman (more frequent, short, continuing, among related processes)
 - Consists of user commands to tools and tool responses to users
 - Some patterns are same as above
 - Often patterns differ: Consecutive requests are related and must be serviced by the same tool
 - Varies from the infrequent, short request pattern to frequent, long transmissions
- d. File package to/from File package (infrequent, very long, among related processes)
 - Small fraction are short, infrequent messages
 - Bulk is files (infrequent transmission of many bits)

5.9 The Cronus Distributed Operating System

5.9.1 Introduction

The Cronus DOS was one of several models for the LAN study, upon which the identification of networking protocols was made. Cronus represents an advanced development which will structurize a new architecture and establish the requirements for interprocess communications. The Cronus DOS plans to employ the DOD's Transmission Control Protocol (TCP) and Internet Protocol (IP) along with a high-speed data bus LAN. The following describes the Cronus DOS and the Interprocess Communications architecture, based on references [29, 30 and 31]. Cronus imposes global resource management, including fault tolerance, and is designed on an object-oriented model.

5.9.2 Cronus, A Distributed Operating System

Distributed systems can be classified along architectural lines according to the physical extent of distribution the system exhibits. We can identify three major architectural areas of interest:

1. Node Architecture (Local)
2. Cluster Architecture (Regional)
3. Inter-Cluster Architecture (Global)

Each of these is related to the emergence in technology of distributed systems, but the technology of distribution tends to be different in the three areas, as explained below.

5.9.2.1 Node Architecture

The development of a processor architecture, configuration, and operating system for a single host or processing node is a large-scale undertaking, usually accomplished by computer manufacturers. A host is typically physically small (can be contained in one room), is designed by computer hardware architects as a single logical unit, and is concerned with maximum event rates of approximately 1 to 1000 million events per second. Although dual-processor nodes have been common for some time, nodes with many-fold internal distribution are just now becoming commercially available. The structure and efficient utilization of such hosts are at the forefront of computer architecture research.

5.9.2.2 Cluster Architecture

A cluster is a collection of nodes (Figure 5.9.2.2) attached to a high-speed local network. At present, technology limits the speed of local networks to approximately 10 to 100 megabits aggregate throughput, and the physical size of the network to a maximum diameter of about 4 kilometers. The host systems are made to work together through the agency of the distributed operating system, which provides unifying services and concepts which are utilized by application software. The maximum event rate at the DOS level is related to the minimum message transmission time between hosts, and is on the order of 10 to 1000 messages per seconds. The cluster configuration and applications supported by it are typically under the administrative control of one organizational entity.

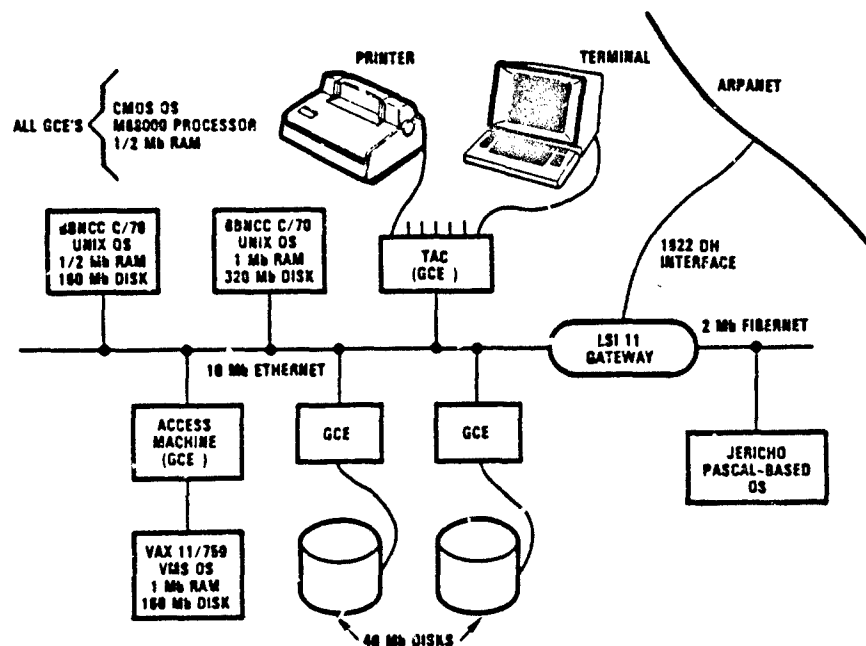
5.9.2.3 Inter-Cluster Architecture

An inter-cluster architecture (Figure 5.9.2.3) typically deals with geographically distributed clusters (or in the degenerate case, hosts) which are not under unified administrative control. Because of administrative issues and the greater lifespan of inter-cluster architectures, they tend to be composed of parts from many different hardware and software technologies. The communication paths between widely separated clusters have much lower bandwidth and higher error rates than local networks; the maximum event rate for cluster-to-cluster interactions is on the order of 0.01 to 10 events per second. In the inter-cluster case, emphasis is on defining protocols for interactions between clusters and on the appropriate rules for the exchange of authority (for access to information and consumption of resources) between clusters.

5.9.3 The Cronus DOS Functions

Expected usage of the DOS can be divided into five categories:

1. Applications
2. Application development and maintenance



10457-12

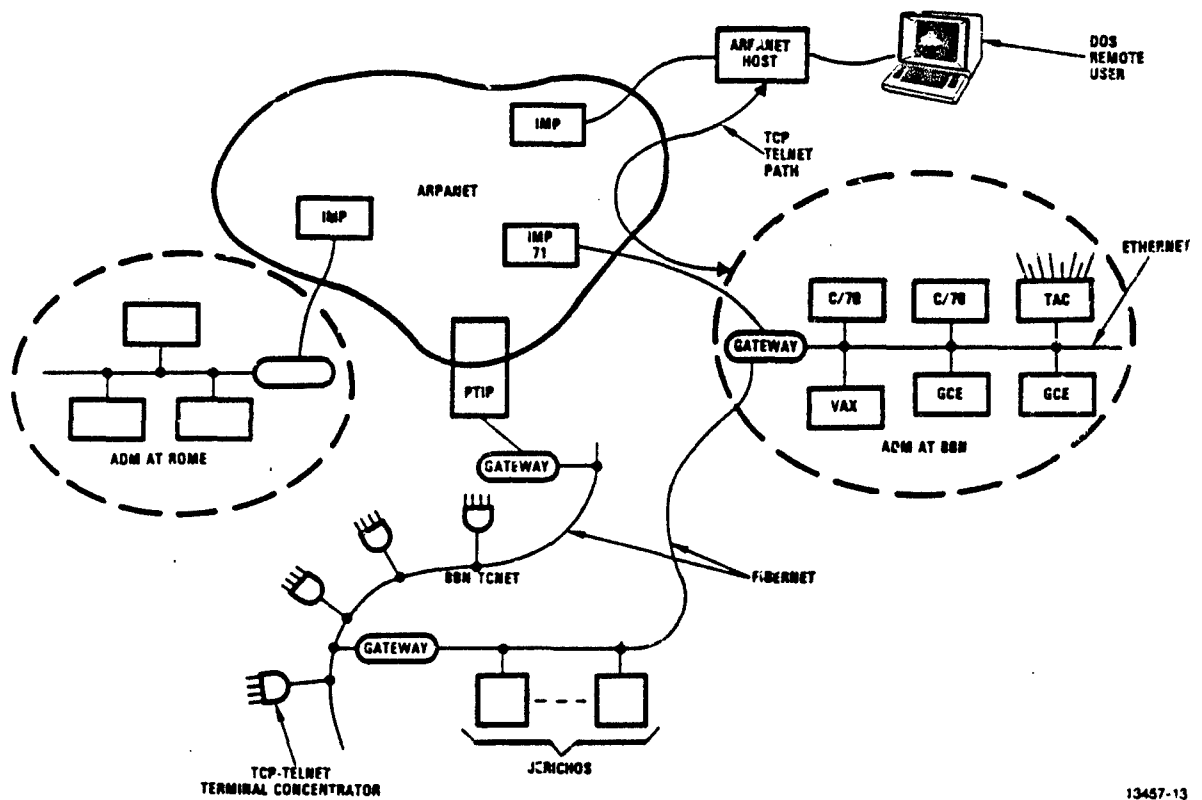
Figure 5.9.2.2. The Local Cronus Cluster Configuration (Physical Model)

3. System administration
4. System operation
5. System development and maintenance

The system is intended primarily to support end application usage. However, to adequately support end applications it must also support the other categories of use.

The DOS system will provide functions in the following areas:

- System Access. The objective is to support flexible, convenient access to the system from a variety of user access points.
- Object Management. The notion of a "DOS object" is central to the user model for the DOS. The DOS treats resources, such as files, programs and devices, as "objects" which it manages, and which users and application programs may access. The objective of the object management mechanisms is to provide users and application program uniform means for accessing DOS objects.



13457-13

Figure 5.9.2.3. The Cronus Inter-Cluster Environment (Physical Model)

- Process Management. The "process" abstraction, also a Cronus object, is central to the user model of the DOS. In addition, it is useful as an organizing paradigm for the internal structure of the DOS. The objective of the DOS process management mechanisms is to implement the "process" notion in a way that enables processes to be used both to support the execution of application programs for users and internally to implement DOS functions.
- Authentication, Access Control, Protection, and Security. The objective is to provide controlled access to the DOS objects.
- Symbolic Naming. DOS users will generally reference objects and services symbolically. Symbolic access to DOS objects will be supported by means of a global symbolic name space for objects.

- Interprocess Communication. The objective of the interprocess communication (IPC) facility is to support communication among processes internal to the DOS, and among user and application level processes.
- User Interface. The user interface functions provide human users with uniform, convenient access to the features and services supported by the DOS resources.
- Input and Output. The objective here is to provide flexible and convenient means for users and programs that act of the behalf of users to make use of devices such as printers, tape drives, etc.
- System Monitoring and Control. The purpose of the system monitoring and control functions is to provide a uniform basis for operating and manually controlling the system.

The principal goal for the DOS in each of these functional areas is to support features that are comparable to those found in modern, conventional, centralized operating systems.

5.9.4 DOS Provides Essential Services System-Wide

At the heart of the DOS concept is the availability of selected, essential services to all of the applications in the DOS. The coherence and uniformity of the DOS is directly enhanced when applications and application host operating systems embrace the DOS-supplied services as the single source of these services. To the extent that applications and application host operating systems choose to utilize parallel but incompatible services, coherence and uniformity is lost.

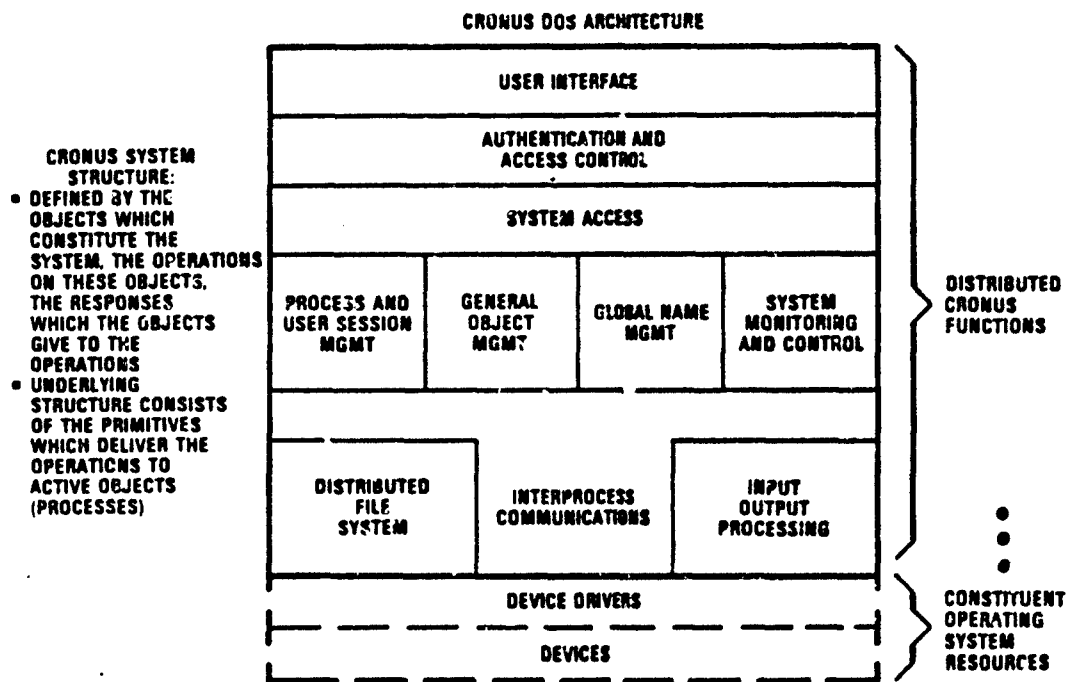
At this time the essential services are:

- User access points (terminal ports, workstations) providing a uniform path to all DOS services and applications
- Object management (cataloging and object manipulation) for many types of DOS objects
- Uniform facilities for process invocation, control, and interprocess communication for application builders
- Cluster-wide user identifiers and user authentication as the basis for uniform access control to DOS resources
- Cluster-wide symbolic name space for all types of DOS objects
- A standard interprocess communication (IPC) facility supporting datagrams and virtual circuits
- A well-designed user interface that provides access to all DOS and applications services

- Input/output services for the exchange of data with people and systems apart from the DOS
- Host monitoring and control services, and additional mechanisms needed for cluster operation

5.9.5 Cronus Logical Model Architecture

While Cronus can be described in its physical sense, this will discuss the LAN Study's logical view of the Cronus architecture. Figure 5.9.5-1 illustrates a view of the Cronus DOS architecture, structured in a top-down view, starting at the top with the user. Some functions at the bottom would be implemented and be under the control of a Constituent Operating System (COS). Those above this level would be implemented on the COS but would be under the control of the global operating system.

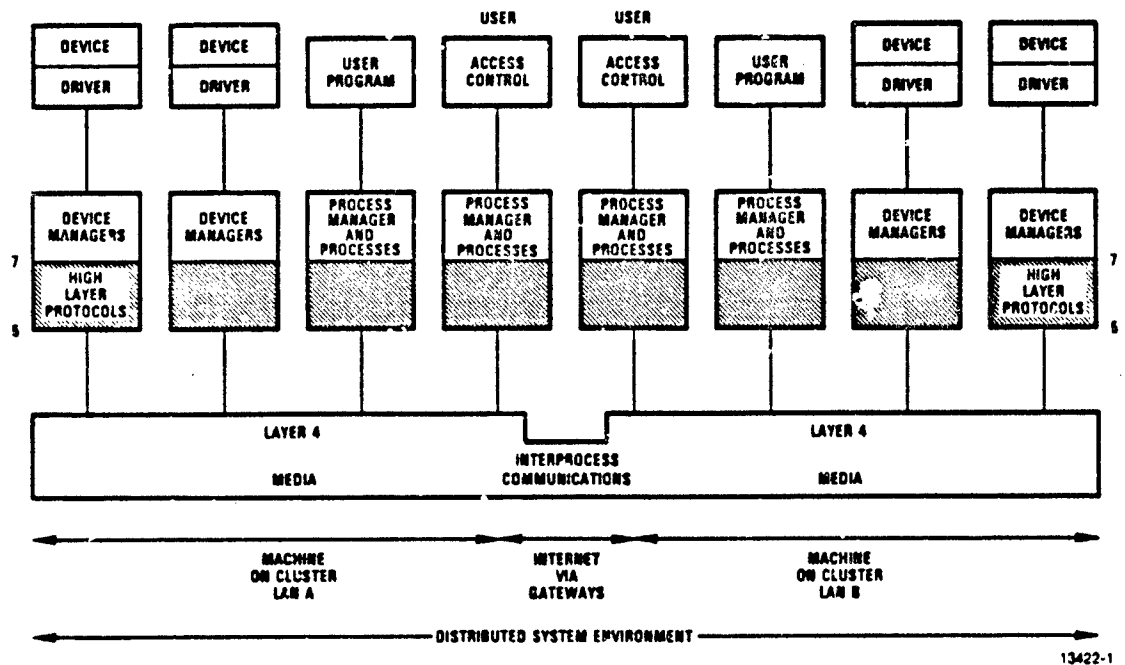


13422-2

Figure 5.9.5-1. CRONUS DOS Architecture (Logical View)

In regard to the Cronus global operating system, it was concluded by the LAN Study Team that in spite of it being called a Distributed Operating System (DOS), the makeup and architecture is that of a Network Operating System (NOS), meeting the criteria given by reference [27].

Cronus is an object-based architecture. Objects can be physical or logical. Objects have object managers. The Cronus objects and their managers can be dispersed throughout a distributed system environment. This is shown by Figure 5.9.5-2. It shows two machines on two different clusters, interconnected by an internet with gateways.



13422-1

Figure 5.9.5-2. Distributed System Environment

Users and objects (devices) interface to their respective object managers. They in turn communicate with other object managers by exchanging data plus control messages. These exchanges follow well-established rules, called protocols. A family, or suite, of peer protocols are needed in order for the Cronus type of resource distribution to properly interoperate. Figure 5.9.5-2 identifies where the LAN Study Team views networking protocols to reside in Cronus. The layers 5-7 set of high layer protocols interface and provide services to the distributed object managers. The high layer protocols correspond to the application, presentation and session layers of the OSI/RM and use the underlying Layer 4-Media Interprocess Communications facility for exchanging Object Manager peer protocol messages. There are tiers of peer protocols considered needed for use between the distributed object or resource manager type entities.

5.9.6 Cronus Cluster Physical Model

The equipment configuration shown in Figure 5.9.2.2 for the DOS cluster is briefly reviewed. The DOS cluster is composed of three types of equipment:

1. Application Hosts. These may be general-purpose hosts (e.g., timesharing machines) providing services to many DOS users, or workstations providing services to one user at a time, or special-purpose hosts (e.g., signal processing computers) required by just one DOS application. Application hosts are often user programmable. In general, they have many characteristics which are not under the control of the DOS; the DOS must be sufficiently flexible to incorporate application hosts of almost any kind.
2. DOS Service Hosts. These machines are dedicated entirely to DOS functions and exist only to provide services to DOS users and applications. In general, they represent modules with specific, system-oriented functions (e.g., archival file storage) and are not user programmable. Requirements for the DOS service host types and operating systems will be specified in the DOS design documents.
3. A Communication Subsystem. The subsystem consists of a high bandwidth, low-latency local network, hardware interfaces between hosts and the local network, device driver software in the host operating systems, and low-level protocol software (the data link layer) in the hosts.

Application hosts will be connected to the communication subsystem in one of two ways: 1) directly, by means of a host-to-local-network device interface; or 2) indirectly, through an intermediary DOS service host called an access machine.

5.9.7 DOS Generic Computer Elements

The concept of a Generic Computing Element (GCE) is important to the DOS design. A GCE is an inexpensive DOS host that can be flexibly configured with small or large memory, and with or without disks and other peripherals, as shown in Figure 5.9.2.2. GCE's will be configured for, and dedicated to, specific DOS service roles, such as terminal multiplexing, file storage, access machines, and DOS catalog maintenance.

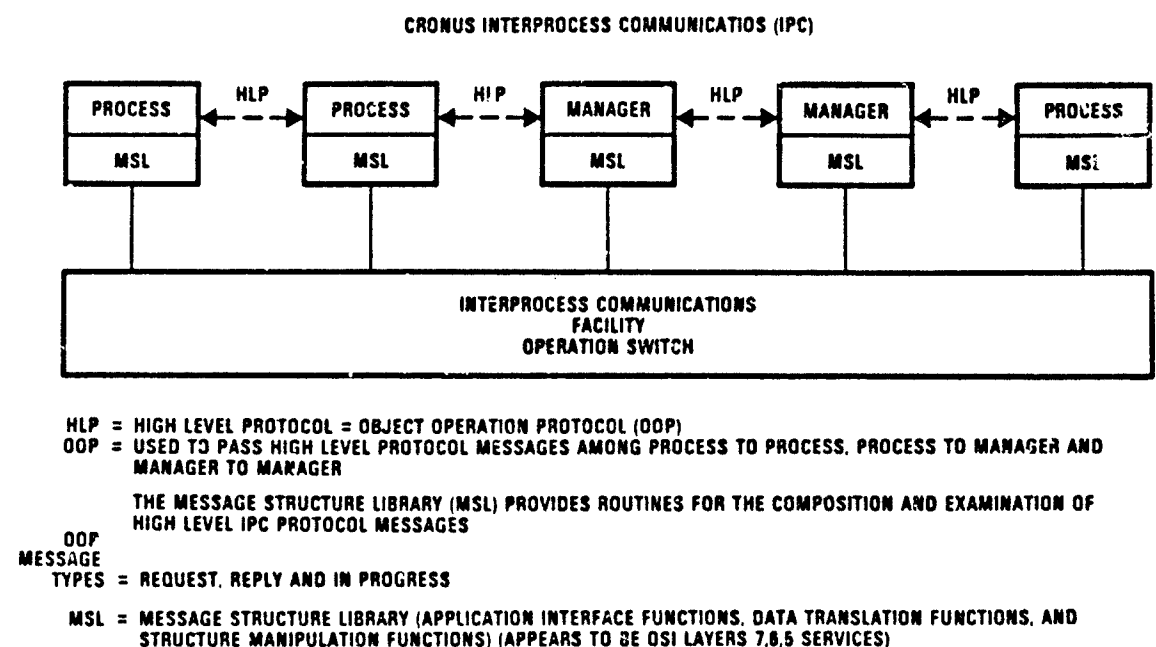
The GCE's are the basis for implementing the essential DOS services in a uniform, application-host-independent manner. Because the DOS design will specify the properties of GCE's and also the software components running on them,

it is possible to control the performance and reliability characteristics of the essential DOS services. A configuration consisting of the local network, some number of GCE's and supporting the essential services represents the minimum useful DOS instance.

Application programs can be constructed above the GCE hardware and operating system; a single GCE host may support DOS services or user applications, but not both.

5.9.8 Interprocess Communication (IPC)

The objective of the DOS interprocess communication (IPC) facility is to support the communication requirements of the DOS. The Cronus IPC is shown in Figure 5.9.8. Requirements can be identified at two levels:



13422-6

Figure 5.9.8. Cronus Interprocess Communication

1. The System Implementation Level. The collection of software modules that implement the DOS. The DOS executes as processes on various DOS hosts. These interactions are supported by the interprocess communication facility.

2. The User Application Level. Some of the application programs that execute in the DOS environment may be structured as distributed programs. A distributed program is one whose components may run as cooperating processes on different hosts. The components of such a distributed application program will need to communicate.

The IPC facilities that are available at the application level will be built upon the system level IPC facility.

The DOS interprocess communication (IPC) facility will be based on the following principles:

- The IPC mechanism will support a variety of communication modes including datagrams and connections (i.e., reliable sequenced, flow controlled data streams).
- It will be built upon the standard DOD IP (internet) and TCP (transmission control) protocols. This assumes that the implementations of the DOD protocols that are used will provide adequate performance (low delay, high throughput). If they do not, it may be necessary to build the IPC directly on the local network (Ethernet) protocol.
- Interhost and intrahost communication will be treated in a uniform fashion at the interface to the IPC facility. That is, the same IPC operations used for communicating with processes on different hosts will be used for communicating with ones on the same host. Of course, to achieve the efficiencies that are possible for local communication, the IPC implementation will treat interhost communication differently from local communication.
- Several levels of addressing will be supported by the IPC facility. The details of IPC addressing within the DOS have not yet been finalized. The fundamental issue which is unresolved is what the addressable entity for the IPC facility shall be; that is, to what will datagrams be addressed and what will connections connect? One reasonable choice would be for the process itself to be the addressable entity. Alternatively, another abstraction, the "port," could be introduced for this purpose. Ports would be objects, and like other objects such as processes, they would have

unique ID's and, if cataloged, could be referenced symbolically by name. Regardless of the choice for addressable entity, the IPC facility will permit addressing by means of unique ID and by means of symbolic name. Other levels of addressing may also be supported. At the interface to the IPC facility wherever IPC address is expected, any of the supported levels of addressing (unique ID, symbolic name) may be used.

- The ability of the IPC facility to deal with symbolic addresses will permit it to support "generic" addressing. This will permit processes to specify interactions with other processes in functional terms.
- The IPC mechanism will provide means to directly utilize some of the capabilities of the local network. For example, the Ethernet supports efficient broadcast and multicast. The IPC will provide relatively direct access to these capabilities by supporting broadcast and multicast addressing. To achieve the design goal of component substitution it is important for the DOS system to be as independent as possible of the specific characteristics of the particular local network chosen for the ADM. Therefore, care must be taken to avoid building dependencies on the particular ADM network technology into lower level DOS mechanisms, such as the IPC. If such dependencies cannot be avoided, care must be taken to minimize their impact on the DOS. In our opinion, this is not an issue in the case of the broadcast and multicast facilities, since many state of the art local network technologies support similar capabilities.

5.9.9 The Communication Subsystem

A high-bandwidth, low-latency local network is the backbone of the DOS. The DOS concept of operation will specify the interface to the local network, so that alternate local network technologies can be substituted for the particular local network chosen for the Advanced Development Model.

The local network will permit every host to communicate with every other host in the DOS cluster and will provide an efficient broadcast service from any host to all hosts. The local network interface specification may further restrict the minimum packet size, addressing mechanism, and other local network properties.

5.9.10 Candidate Protocols for Cronus

The Cronus DOS design employs several protocols. Some, like the underlying Ethernet and Transmission Control and Internet Protocols (TCP/IP), are industry and DOD standard forms. These provide the basic interprocess communications. Others, like the MSF, SER, MSL, OS, and OOP, are new ones developed for Cronus. These perform higher layer type protocol functions. They correspond to the OSI/RM layers 5, 6, and 7 in the following ways:

- a. Control and Stream message-based protocols
- b. Cronus Message Structure Facility (MSF) consisting of a Message Structure Library (MSL) and a Standard External Representation (SER)
 - The MSL seems to be the same services/functions which OSI Presentation Layer is to perform.
 - The SER deals with data structuring, also part of OSI Presentation Layer.
- c. Operation Switch (OS)
 - The OS functions dealing with asynchrony and demultiplexing correspond to the OSI Session Layer.
- d. Object Operation Protocol (OOP)
 - The OOP appears to correspond with some of the OSI Application Layer protocols for object-based distributed processing

Close coordination between Cronus and OSI protocols might permit employment in the future of higher layer industry protocols as that work matures and could be fitted into and upgraded or productized form of Cronus. The following subsection discusses a Generic Network Operating System, called GNOS, which is an approach to avoiding the probable heterogeneity the current Cronus exhibits.

5.10 Generic Network Operating System (GNOS)

5.10.1 Introduction

To avoid the development of a closed or heterogeneous global operating system set of networking protocols, a general representation of a Network Operating System, called GNOS, was developed. With the emerging international OSI attempt to build a family of open system interconnection protocols to create a

global open networking system, the LAN Study Team took the view that a generic, rather than a possible proprietary form such as the Cronus global operating system, formed a better base to construct command, control and communications protocols upon. This subsection discusses GNOS, its architecture, subsystems, services, functions and protocols.

5.10.2 Objectives for GNOS

The following is a statement of objectives and purpose for GNOS:

Objectives

- To describe a generic visualization of a distributed processing Network Operation System (NOS) which exhibits the capabilities to manage the resources of a collection of connected computers, peripherals, input/output devices and a means for interprocess communications. GNOS defines the functions and interfaces available to application programs on system hosts.
- GNOS specification provides an open* reference model to focus protocol standardization, study and development works upon.
- GNOS provides the Air Force a reference model to guide development of C³ technology, protocols and systems.

Purpose

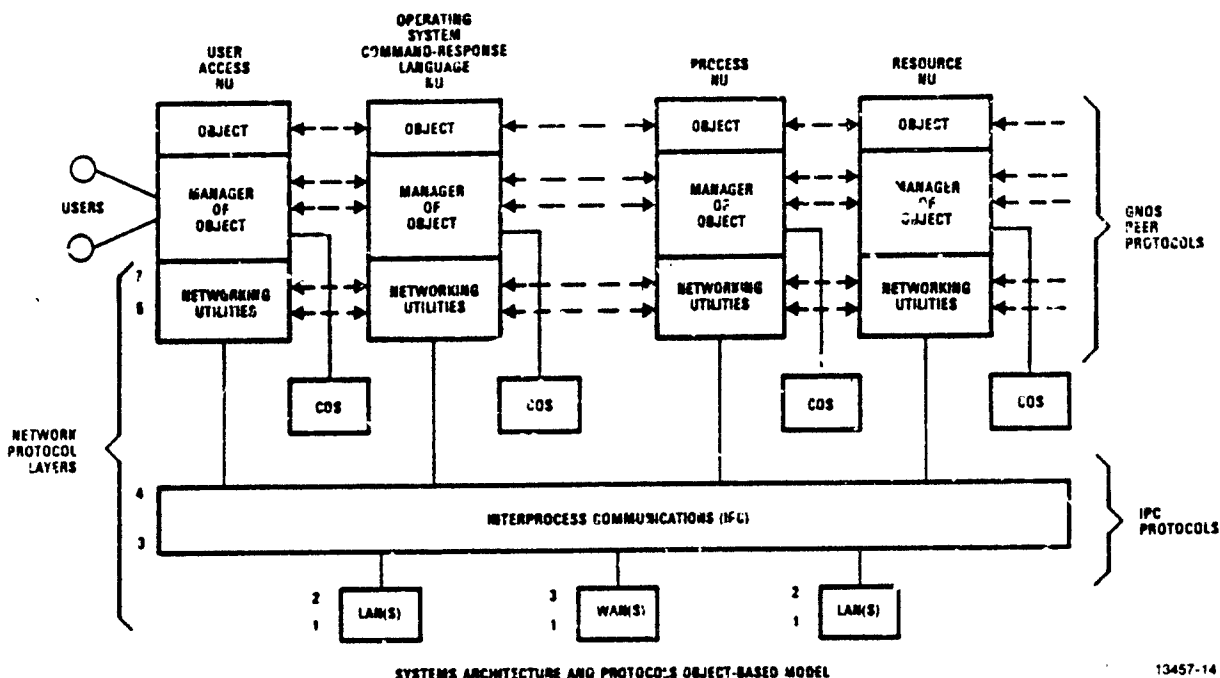
- Provide a unified representation of an integrated distributed processing network operating system in terms of its architecture, subsystems, services, functions and protocols.

5.10.3 GNOS Architecture

The GNOS Architecture consists of the following principles and is illustrated by Figure 5.10.3:

- Object oriented abstract design.
- System consists of a collection of objects that are accessed or updated by users through transactions.
- Users access applications programs, which access object (resource) managers. These access local resources by way of a Constituent Operating System or remote resources by way of networking utilities protocols and interprocess communications protocols. Presentations Layer provides unifying information representation and interface for the resource managers to the interprocess communications and data storage/retrieval subsystems.

*Open means it is vendor independent.



13457-14

Figure 5.10.3. Generic Network Operating System (GNOS)

- Parallel processing programming language exhibiting strong data typing implements application programs.
- An "Operating System Command and Response Language (OSCR)" provides a global system-wide command language for users and processes to access system resources.
- Resource manager's protocols and networking utility protocols are grouped as a unified subsystem of interconnected Network Units (NU's).
- Network Units cooperate by passing messages to each other via interprocess communications.
- The interprocess communications supports intra-host and local/wide area inter-host services via a loosely coupled suite of message passing networking protocols.
- Distributable functions supported:
 - Application programs
 - Application subsystems
 - Processing

- Data Base
- Communications
- End-user devices

5.10.4 GNOS Subsystems

The following subsystems were identified for GNOS:

- User terminals and terminal managers.
- General purpose processors and peripherals.
- Storage (short-term, long-term).
- Interprocess communications.
- Constituent Operating System(s)
- Network Operating System comprised of resource managers and networking utilities grouped into Network Unit subsystems (NU's).
- Application programs and subsystems.
- Processing environment.
- File and Data Base.

5.10.5 GNOS Services

The following were identified as services which GNOS provides:

- User access/authentication
- Multilevel security
- Object management
- Process management
- Resource management
- Networking utilities (remote access to files, terminals, jobs, messages)
- Operating System Command and Response Language (OSCR)
- Task assignment and management
- Directory (symbolic name to network identifier)
- Data base management
- Input/output
- Monitoring and control
- Supports transaction and job processing
- Program generation, invocation and support
- Access local resources via constituent operating system interface
- Resources managed:
 - Processor cycles
 - Main storage
 - Files

- I/O devices
- Sessions, queues, data base records

5.10.6 GNOS Functions

The following GNOS functions were identified:

- Resource allocation and management (programs request access to resource via a request to its object manager. Higher functions may come into play to recover, secure, change control); these functions can be system-wide, network remote or local on its Constituent Operating System.
- Provides user interface, authentication/access protection and system access.
- Common command language interpreter.
- Interprocess communications (connection-oriented for streams and connection-less for transaction messages).
- Process management, distributed task scheduling, multitasking and concurrent processing performed.
- File access, transfer and manipulation.
- Job transfer and management.
- Distributed data base access, update and management.
- System monitoring and control, fault tolerance and system recovery.
- Provide access to local operating system functions.
- Directory naming/addressing.
- Capacity management.
- Configuration and reconfiguration management of processors, cluster, global system.
- Support information processing of data (creation, distribution, storage, manipulation, output).
- Multilevel security (authentication, resource access, encryption).

5.10.7 Protocols Needed to Support GNOS

This subparagraph discusses the protocols identified as being required to support the GNOS. Message-based transaction and file transfer protocols, in support of software program functions, comprise the Generic Network-wide Operating System (GNOS). Software program functions perform distributed resources (object) management operations in support of the distributed applications (policy setting/execution functions reside here) by way of resource manager protocols and assessing networking utility services.

A networking protocol suite is comprised of three service regions:

- Networking-Wide Utilities (Canonical form of high level services for accessing remote resources)
- Host to Host/Internet data transport services
- Local/Wide Area Network transmission services

Networking-wide Utilities provide generic services to the resource (object) manager entities to enable remote resource access to files, terminals, data, jobs, messages, etc. The combination of resource (object) manager and networking-wide utilities functions comprise the distributed Network Operating System (NOS); where a local resource is employed, the resources manager accesses it using its Constituent Operating System (COS).

Resource Managers make a set of resources available to users (programs), such as processor cycles, main storage, files on disk or tape, such I/O devices as keyboards or displays, and such abstract resources as sessions, queues, or data base records. Resource Managers allocate resources to users as its central function in response to a user request. This includes access scheduling, coordination, resource allocation deadlock detection, resource change commitment control, resource access security and resource formatting services. Resource Managers employ resource coordination peer protocols and access network services by way of interfacing to the networking-wide utility protocol services.

Program to Program protocols exchanged between Resource Manager/Network Utility entities make up the distributed Network Operating System. Protocols support cooperation among the distributed resources managers to perform the following activities:

- Interprocess communications
- Data representation
- Data storage (media, file and data base)
- Process management
- Resource management
- Integrity and security
- Program support

Protocols are needed at several levels to enable distributed object (resource) managers to cooperate autonomously. The levels identified consist of the following:

- User to OSCRL
- OSCRL to Resource Managers
- Resource Managers to Resource Managers

- Networking Utilities to Networking Utilities
- Host to Host interprocess communications
 - Transport
 - Internetwork
- Subnet transmission
 - Local area network
 - Wide area network

Networking wide utility protocols required are*:

- Network management
- Virtual Terminal
- File access, transfer, management
- Job transfer/manipulation
- Message handling
- Document interchange
- Name server
- Data base access/management

Gateways are required for interoperability

- Intra-system and intersystem connectivity functions performed

The following protocol characteristics were defined:

Interprocess communications facility exhibits a loosely coupled message passing characteristic (not memory sharing). Transactions to be the primary method for objects to communicate through exchanging messages; however a connection-oriented service to be required for transferring files.

Message exchange characteristics:

- Intra-host (local only)
- Inter-host (LAN only, LAN/internet/LAN)

Message types:

- Small, minimal effort (datagram service)
- Small, reliable (virtual circuit service)
- Large, reliable (virtual circuit service)

Predominant form of messages exchanged will be for control. These request operations to be performed on objects (local/remote), with replies generated by performed operations, plus exception notices and messages to coordinate the distributed object managers.

*Supported by its type presentation and session protocols

- Standardized Object Manager to Object Manager message types and structuring is required of message formats employed.
- Asynchronous handling of simultaneous process to process transaction messages is required
- Stream interprocess communications required between cooperating processes; has a uni-directional data channel session type between two objects; one is a data source, the other a data sink; connects processes with files, devices, and other processes.
- Object Manager to Object Manager communications employs a high level form of protocol, is asynchronous and involves handling interleaved messages from possibly several processes. Messages are received, requests are originated to satisfy the client requests, and a reply message sent to the original message. In the case of a failure, the object manager assures the client that either all changes requested will take place, or none will for the atomic transaction performed.
- The following resource (object) Managers require peer protocols:
 - Program
 - Terminal Manager
 - Authorization/access control/security
 - OSCRL
 - Catalog
 - File
 - Device I/O
 - Monitoring and control
 - Network Management
 - Data base
 - Directory
 - Host

5.11 Comparison of DOD and ISO Networking Protocol Reference Models

Both the DOD and ISO reference models have significant impacts upon the world of networking. Therefore, it is important to gain insight into their similarities and differences. This subsection provides the results of examining these issues. The following paragraphs provide expanded discussions on both of these.

5.11.1 Layered Architectures

Before comparing the architectures of the International Standards Organization's Open Systems Interconnect model (ISO OSI) to the Department of Defense's model it may be well to revisit the principles of protocol layering. Both of the models are based upon this concept.

A layered protocol architecture provides a hierarchy of control by functionally decomposing overall network communication objectives into strata. Each stratum, or protocol layer, is supposed to provide a particular set of services. Starting at the lowest layer, the services of each succeeding layer are available to the layers above and are built upon the layers beneath. The lowest layers are more physical, the upper ones are more virtual in their makeup.

The advantages of this approach are that it allows for local optimization (important because of the heterogeneous nature of today's systems) while preserving the ability to establish common communications conventions. There is also the hope that this layered approach will minimize the impact of technical evolution by allowing functional replacement of a layer. The practicality of this latter point and, in fact, its necessity can be questioned.

The functionality of a layer in no way implies an implementation scheme. Layered architecture allows the protocol designer the freedom to implement the function according to the particular environment and requirements.

5.11.1.1 Formal Versus Soft Layering of Protocols

Ever since the architectural design concept of layering was introduced, particularly evident in the network architectures of the ARPANET, OSI/RM and proprietary ones like IBM's SNA, there has been concern expressed over the implied complexity or overhead penalties incurred. This paragraph discusses these issues based upon References 5, 10, 14, 19, 20, 25, 40, 42, 43, 44, 60, 82, and 92.

Layering is a design process whereby a set of very complex and interrelated functions are grouped in a common way to form a set of hierarchical elements, called layers. This process breaks an otherwise complex problem into an ordered set of manageable pieces. The grouping of functions into a layer is done in such a way as to find the combination which produces the minimum set of interface operations necessary between two adjacent layers. A layer can also be thought of as a module.

Layers of equal functionality can be distributed across physical machines. When this is done in a networking environment, the layers cooperate with equal, or peer, layers by exchanging data and control information. This is done through message exchanges. These exchanges must follow a prescribed set of

rules which are called peer protocols. The full set of peer protocols comprise the networking communications protocols.

Reference [42] discusses the advantages and disadvantage of protocol layering and presents the points given in Table 5.11.1.1. In summary, it stated, "In general, the advantages are great, the disadvantages slight."

5.11.1.1.1 Layering of the OSI/RM

In the OSI/RM - Communication takes place between application processes running in distinct systems in which a system is considered to be one or more autonomous computers and their associated software, peripherals, and users that are capable of information processing and/or transfer.

In the OSI/RM Architecture - All services provided to the application process users are grouped into seven subsystems of protocol layers, whose entities are distributed among the interconnected systems and which communicate by message exchange among entities of equal layer in the structure according to formal rules of communications called peer protocols. This layering technique is similar to the one used in structured programming, where only functions performed by a module (and not its internal functioning) are known by its users.

Layer Services - These are the capabilities which a layer provides to a user of that layer. Within a layer, functions will be performed, some of which are not user provided services. Only capabilities seen by the user are the services.

A Service Specification - Represents a lower level of abstraction that defines the service provided by each layer. Tighter constraints on the protocol and the implementation that will satisfy the requirements are given. It defines the facilities given to the user and an abstract interface for each protocol layer by specifying the primitives the user would invoke to access the service.

Layer Service Access Points - Abstract ports through which users request and receive the services provided by a particular layer.

Layered Functions - These are the activities which are performed within a layer. Functions are performed by layer entities through cooperation with other layer entities distributed throughout the network. Entities of equal layer value communicate by peer protocols.

Protocol Specifications - Represent the lowest level of abstraction in the OSI standards scheme. Each protocol specification defines precisely what control information is to be sent and what procedures are to be used to interpret

Table 5.11.1.1. Advantages and Disadvantages of Layering

Advantages of Layered Architectures:

1. Any given layer can be modified or upgraded without affecting the other layers.
2. Modularization by means of layering simplifies the overall design.
3. Different layers can be assigned to different standards committees, or different design teams.
4. Fundamentally different mechanisms may be substituted without affecting more than one layer (e.g., packet switching versus leased-line concentrators).
5. Different machines may plug in at different levels.
5. The relationships between the different control functions can be better understood when they are split into layers. This is especially true with the control actions which occur sequentially in time from layer to layer.
7. Common lower level services may be shared by different higher level users.
8. Functions, especially at the lower layers, may be removed from software and built into hardware or microcode.

Disadvantages of Layered Architectures:

1. The total overhead is somewhat higher.
2. The communicating machines may have to use certain functions which they could do without.
3. To make each layer usable by itself there is some small duplication of function between the layers.
4. As technology changes (e.g., as cryptography and compaction chips become available, or these functions can be built onto HDLC chips) the functions may not be in the most cost-effective layer.

In general, the advantages are great, the disadvantages slight.

this control information. The protocol specifications constrain implementations sufficiently to allow open systems to communicate while still allowing differences in implementation.

5.11.1.1.2 Layering in the DOD/RM

The ARPANET/DOD architecture and protocols suites for networking predate the work of the ISO in developing the open systems interconnection architecture and protocols. The Government now has two networks, split physically; ARPANET (for the R&D community) and MILNET (for DOD operational users). Layers of protocol functionality between the DOD and the ISO agree well up through the transport layer (this is based on both supporting at least connection-oriented, connectionless and internetworking protocols). Layers of protocol functionality above the transport provide similar services but are very different in the packaging of functions. The DOD's approach has been to group service functions/protocols more into subsystems whereas the OSI has maintained clean separation of the layers for session, presentation and application. This has resulted in a vertical packaging of functions by DOD (with some intermixing of functions across layers 5, 6 and 7) and a horizontal packaging into layers by OSI, with some added interface formality between layers.

This vertical packaging of protocol layered functions into somewhat of a subsystem is very similar to the way IBM's SNA has done its higher layer protocols. In SNA [82], the Transport through Applications layer protocols are packaged vertically into what is called a Logical Unit, or LU. The LU is made up of subelements of layers 4-7 protocols and are configurable when a session is set up. LU's are given Type designations (i.e., LU 0-6.2) which delineates end user properties. The DOD's approach, based upon experience emanating from the ARPANET research community, has characteristics similar to, but not as extensive or formal as, the SNA Logical Unit.

5.11.1.1.3 Soft Layering of Protocols

Cooper [43] examined the issue of soft layering of protocols as part of his masters thesis at MIT in 1983. He examined the issue of the efficiency of protocol layering. He found two major sources of inefficiency in protocol implementation. These were caused by "the imposition on them of a layered structure."

In the thesis abstract the following is stated:

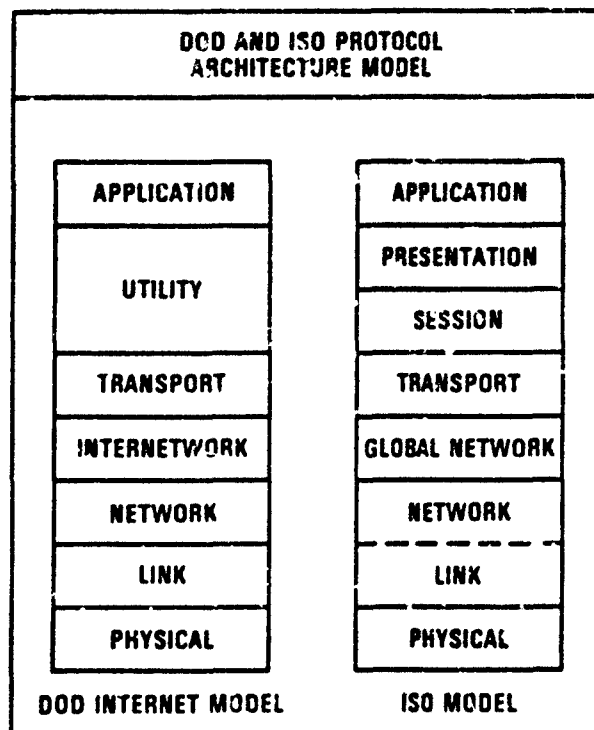
"The conventional approach to making layered protocol implementations run efficiently - for avoiding the sources of inefficiency - are all

independent of the protocol specification, and thus all decrease the value of the protocol specification as a guide for implementing protocols."

The report "introduces a new means of avoiding the problems of layered protocol implementations which operates within the domain of the protocol specification. We allow an increase in the flow of state information between the layers of a layered protocol implementation in a very controlled manner so as to decrease the modularity of the protocol architecture as little as possible. Since our approach decreases the rigidity of the layered structure without entirely eliminating it, we coin the term "Soft Layering" for the approach."

5.11.2 DOD Versus ISO Models

Figure 5.11.2 is a simplistic representation of the DOD and OSI layered protocol architectural models. Organizations which participate in the ISO networking architecture/protocols standards making work are numerous throughout the world. Table 5.11.2 lists only a few. These are CCITT, ECMA, ANSI, NBS, IEEE, EIA, DOD, special interest groups and industrial organizations. On the other hand, DOD is more of a closed organization without the participation by such a cross section as in the ISO.



13277-23

Figure 5.11.2. DOD and ISO Reference Models

Table 5.11.2. Groups Involved in Standards

Groups involved in standards			
Organization and geography	Affiliation	Membership and representation	Influence
ISO (International Standards Organization) International	Voluntary, nontreaty	Standards bodies in participating nations. U.S. representative is ANSI; ECMA is observer	Responsible for Open Systems Interconnection model. Close relationship with CCITT
CCITT (International Consultative Committee on Telegraphy and Telephony) International	Part of International Telecommunications Union (a U.N. treaty organization)	Private companies; scientific and trade associations; postal, telephone, and telegraph administrations. U.S. representative is Department of State	"Recommendations" which are law where communications in Europe are nationalized
ECMA (European Computer Manufacturers Association) Western Europe	Computer suppliers selling in Europe; includes some U.S. companies	Trade organization of suppliers; small, with about 20 members	Contributes to ISO and also issues own standards known for fast movement
ANSI (American National Standards Institute) United States	Voluntary	Manufacturers, organizations, users, and communications carriers	U.S. voice in ISO
NBS (National Bureau of Standards) United States	Government Agency	Government agencies and network users; much work done by Bolt, Beranek & Newman, which is largely responsible for DOD's Arpanet	Issues Federal Information Processing standards for equipment sold to federal government. Department of Defense need not comply
IEEE (Institute of Electrical and Electronic Engineers) International	Professional Society	Dues-paying individuals	Contributes to ANSI and issues own standards such as IEEE-802.3 CSMA/CD and 802.4 token-bus LANs and IEEE-488 bus
EIA (Electronic Industries Association) United States	U.S. trade organization	Manufacturers	Contributes to ANSI, known for physical layer's RS-232-C standard
DOD (Department of Defense) United States	Government Agency	Government/military	All customers dealing with the military establishment
Special interest industry groups, such as the ANSI X9 banking standardization group	Voluntary organizations, such as ANSI and IEEE	Organizations and firms with a specialized interest	Issues standards that meet a specialized need
Industrial organizations	Self	-----	Determined by market impact

13457-15

An immediate observation is that the functional decomposition of both is the same through the transport layer. It is only at the higher layers that differences appear.

The reason for this could be that the functional basis for the decomposition changes at this point. Through the transport layer all layer objectives are concerned solely with the transportation of data. A protocol above this layer need not be concerned with physical transmission, routing, involvement or existence of intervening nodes, or errors detected and recoveries made. Instead, above the transport layer, the objective is the accomplishment of a particular task. Both the ISO and DOD models concern themselves with this overall goal, namely, the ability to achieve a common (distributed) task.

Recognition of a functional transition above the transport layer is important. It forces a change of perspective. The higher layer protocols look downward to the transport and lower layers only for the data transmission services necessary for the accomplishment of their distributed tasks. At these higher layers there is less of a stratification of functions. Instead (and particularly as viewed in the DOD model), there appear to be groupings of decidedly different functions within the same layer. This is because of the common need for data transmission services but the differing purposes of tasks.

An example of this type of functional transition can be found in our postal system. Here the mailbox acts as the interface between the data transportation layers and the task oriented layers. Contents of mailed letters can be viewed as task data. A person mailing a letter is not concerned with the method of transportation, be it truck, train, or plane. The concern is only that the letter arrive at its destination in some reasonable length of time. This is the service upon which the mailer depends. In turn, the postal department has no knowledge of the contents of letters ("tasks"), be they bills, invitations or correspondence.

5.11.2.1 Application Layer Differences Between ISO and DOD

ISO/RM: Provides a superset of DOD services and identifies management services/functions more explicitly than does DOD for application-processes and OSI system resources. ISO adheres to formal layering structure.

DOD/RM: Some services, such as establishing authority to communicate and privacy mechanisms, have been partially moved to the session layer. As DOD requirements mature, greater divergence from ISO is expected. Management services/functions for applications-processes and DOD resources are not explicitly specified. DOD applies protocol hierarchy structure rather than a strict layering structure.

5.11.2.2 Presentation Layer Differences Between ISO and DOD

ISO/RM

1. Views a resource strictly as a data structure, with a set of commands to access or modify the data; this, according to DOD, is too strict a view.
2. The particular transfer syntax is negotiated at the time of connection establishment, whereas in DOD a default syntax is defined which all must implement, plus a negotiating capability.
3. Quality of service is absent in ISO but is explicitly addressed by DOD.

NOTE: ISO model stresses point-to-point two-party interactions, whereas DOD supports multi-entries distributed.

5.11.2.3 Session Layer Differences Between ISO and DOD

ISO/RM

The ISO session layer's functions are seen by DOD as not adding very much value to the underlying transport layer services. In addition, voice and other real-time stream applications are not adequately supported. ISO only supports point-to-point communications between using entities and does not support connectionless service.

DOD/RM

The DOD session layer differs in four ways from ISO:

1. Support for real time quality of service using entities, for diverse applications
2. Support for distributed applications
3. Provisions for communication between more than two entities
4. Data quantitative has been dropped

In the future, DOD anticipates needing some access control and name service functions added, as well as increased robustness and control.

5.11.2.4 Transport Layer Differences Between ISO and DOD

ISO/RM

ISO transport layer protocols are currently connection-oriented. ISO has less negotiable quality of service. Session layer services are separate from transport.

DOD/RM

A major difference is that DOD stresses use of connectionless (datagram) UDP service as well as connection service TCP at the transport layer. DOD provides greater ability to negotiate quality of service in order to support

wide range of usage (e.g., transaction, bulk transfer and real time). Some session layer services have been merged into transport layer services in DOD.

5.11.2.5 Internet Layer Differences Between ISO and DOD

ISO/RM

There is no ISO internet layer but X.75 is a connection-oriented virtual circuit approach residing in the network layer.

DOD/RM

The DOD approach is different from ISO fundamentally. ISO's approach is based on X.75 type connection-oriented tandem interconnecting of subnetworks which results in virtual circuits serially connected and mapping done between them. DOD, rather, is based on use on connectionless datagram, spanning end-to-end through gateways which join different method subnets together. DOD uses the Internet Protocol (IP) as its internetworking standard. This is more robust, survivable and flexible than virtual circuits in tandem. DOD does not preclude using connection-oriented internet protocols

5.11.2.6 Network Layer Differences Between ISO and DOD

ISO/RM

The view is end-to-end tandem relays of individual subnetworks, with emphasis (currently) on connection-oriented services with X.25 packet switched services as the primary view.

Internetworking is viewed to be done in an upper sublayer of the network layer.

DOD/RM

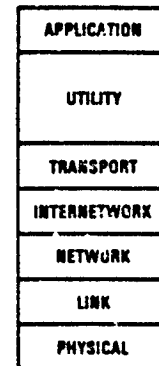
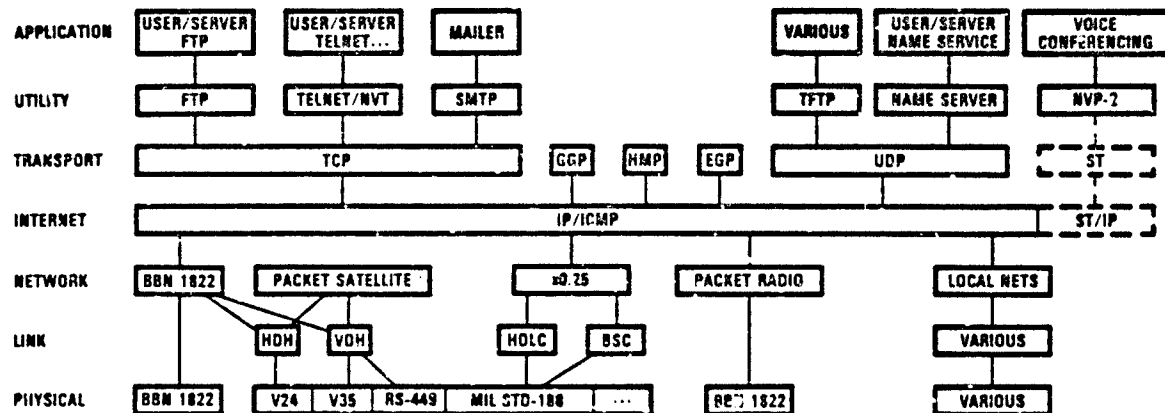
The view is that of services provided by a single network, operating as one of several within an internet, with emphasis on connectionless (datagram) Internet Protocol users. DOD predicts that connectionless (datagram) local area networks will predominate. Internetworking is done in a separate layer above the network layer.

5.11.2.7 Data Link and Physical Layers

There is essentially no difference between the DOD and ISO reference models for the Data Link and Physical layers in the services provided. The ISO has developed for the most part with a connection-oriented perspective and almost exclusive use of carrier based network facilities. The DOD model supports both connection and connectionless services and recognizes more than just public network facilities. The DOD recognizes local area network facilities whereas the ISO model is just beginning to consider the use of LAN's. ISO currently is reviewing the IEEE 802 LAN standards for international use.

5.12 DOD Networking Reference Model

The DOD Protocol Reference model, shown in Figure 5.12 and Table 5.12, is intended to document the principles of protocol design. It is to be a baseline serving the development of standard DOD protocols. It attempts to describe ARPANET and Internet programs and commercial world design principles and experiences. It attempts to prescribe principles for development of future DOD protocols. The DOD/RM diverges from the OSI/RM of ISO in that:



DOD INTERNET MODEL

13277-24

Figure 5.12. DOD Internet Protocol Hierarchy

1. DOD has specific communications requirements (e.g., security, survivability).
2. OSI/RM is too restrictive on the designer.

It communicates basic principles of the "DOD approach to protocol design" to interested parties. Lastly, it is proposed as a guide to the basic direction of DOD protocol specification efforts.

Table 5.12. DOD Reference Model Layers

- Application Layer - Protocols provide distributed information services to an application, to its management, and to system management
- Presentation Layer - Protocols provide virtualization of data formats and distributed resources
- Session Layer - Protocols help coordinate use of multiple transport services, provides name services and access controllers
- Transport Layer - Protocols provide process-to-process communication across one or more networks (Host-to-Host)
- Internet Layer - Protocols perform network to network routing
- Network Layer - Protocols that are network-specific that allow data transfers over a single network
- Data Link Layer - Protocols manage data transfer across a single data link
- Physical Layer - Protocols that manage access and data transfer with a physical communications channel

The general strategy followed in developing the DOD Protocol Reference Model was to use the ISO Reference Model (RM) for Open Systems Interconnection (OSI) as a base; a second major influence was the ARPA Internet project. Where the OSI baseline was thought to be inadequate, the DOD/RM reflects modifications made to meet particular DOD requirements. The DOD/RM report discusses five specific areas that affected the DOD/RM.

- Anticipated DOD network applications
- Internetworking with present and planned DOD (and non-DOD) systems
- Security requirements
- Robustness and other DOD quality-of-service issues
- Phased evolution from existing DOD systems and protocols

Since the DOD/RM (and many of the associated protocols) is not completely specified at this time, it can only serve as a guide. Certain protocols that are already DOD standards are included in DOD/RM. Other protocols, while not yet DOD standards, are actively under development, and designers should be aware of these efforts. In areas where no specific DOD work is under way, the work of national and international standards groups should be monitored.

Regardless of the particular protocols employed, the principles of layering imposed by the DOD/RM (and the OSI RM) should be scrupulously observed. The interfaces between layers should be independent and modular in order to allow

protocols at the various layers to be replaced or extended when it is advantageous to do so (e.g., to accommodate new applications). Correspondingly, the protocols employed in the LAN architecture should follow DOD, national, and international standards. That is, only vendor independent (nonproprietary) protocols should be considered an acceptable approach.

The ARPANET/DOD architecture and protocols suites for networking predates the work of the ISO in developing the open systems interconnection architecture and protocols. The Government now has two networks, split physically ARPANET (for the R&D community) and MILNET (for DOD operational users).

Applications Layer

This provides the means for application programs to access the DOD internetworking environment serving as the windows between communicating application processes. The protocols define information transfer semantics services. The evolution of a multihost, multivendor LAN-based command center environment will likely require a Common Command Language Protocol that provides a common user interface to network functions. Additionally, NBS work in the areas of Network Interprocess Communication and Distributed Data Protocols should be monitored for possible later inclusion in a LAN Protocol Architecture.

Utility Layer (Presentations/Session)

Performs the virtualization of data and resources, preserving meanings while resolving syntax differences. This allows users to transparently manipulate remote resources. DOD places the following services in the utility layer(s):

- Virtual terminal
- Virtual file store
- Distributed data base
- Teleconferencing
- Advanced messaging

It is expected that the above protocols will provide the basis for the presentation layer. In addition, work is proceeding within ISO and NBS in the areas of Virtual Terminal Services, Job Transfer and Manipulation Services, and Management Protocols that should be monitored for possible inclusion in the LAN-based Protocol Architecture.

The inclusion of a Virtual Terminal Protocol (VTP) in the LAN Protocol Architecture is of particular interest. A VTP will enable new terminals to be integrated quickly into the LAN without requiring application modifications; only the interface for a particular terminal type need be aware of its specific characteristics. The objective of a VTP is to accommodate a wide variety of

terminals with capabilities varying from simple character-oriented transmission to advanced capabilities such as color, support for multidimensional data structures, and graphics (e.g., raster, vector, or bit-map). While there has been considerable research on VTP's, for example, requiring a VTP for the LAN will entail significant developmental risks, due to its complexity and the lack of implementation experience in this area.

The DOD reference model identifies some important deficiencies in the ISO session layer services: lack of support for nonrecord-oriented transaction exchanges (e.g., voice, real-time data streams); lack of quality-of-service negotiation between the session layer and its users' inability to coordinate activities of more than two presentation layer entities (e.g., conferencing). The DOD/RM describes a multipoint notion of "session" (different from that of the ISO session layer) established by a presentation layer entity called the "session-controller." It is anticipated that several different session layer protocols might be designed and implemented depending on the applications involved.

Currently, no DOD standard specification for a session layer protocol has been completed. Work is progressing on a DOD Session Control Protocol (SCP), however. The session layer functions required by existing presentation layer protocols are provided by TCP. In some of these presentation layer protocols, additional functions usually associated with the session layer are handled by the particular presentation layer protocol. As a result, a separate session layer protocol is not initially required. Nonetheless, the LAN protocol design should not preclude fully distinct transport, session, and presentation layer independent protocols.

TELNET - Virtual Terminal Protocol

This is a byte-oriented communications facility which uses a TCP connection for:

- Terminal to process
- Terminal to terminal
- Process to process

Based on a scroll mode ASCII TTY, the protocol defines options negotiable between using parties. TELNET does not provide capabilities for properly using newer bit mapped displays where windowing and font selection are used. One approach to improving TELNET is to create a new class of virtual terminal for the latest terminal technology and expansion for future growth.

File Transfer Protocol (FTP)

Promotes file sharing and use of remote computers while shielding the user from variation in local file storage systems. TELNET uses two TCP connections; one for TELNET to set up conditions of control for initiating data transfer, and the second for exchanging the file data across.

FTP transmission modes are stream, block, and compressed. Limitations exist with handling number representation and word length between processors, especially with floating point number representation. General file transfers between widely diverse equipment is limited. In handling specialized, extremely large volumes of file data transfers, development of new protocols might be warranted. Also, the protocol should relieve the user of having to supply many of today's functions by including type information within the block transfer mode header.

Internet Name Server

This uses the user datagram protocol by returning an internet address message in response to a user initiated name request message. As networks for distributed processing grow, a more general distributed system of name servers and their protocols needs development. This needs to handle multiple internet subnetworks, nodes and those that will be mobile.

Transport Layer

Provides network-transparent data transfer between utility layer users. These perform end-to-end functions between originating and destination hosts. Two major types of service are provided; connection-oriented (for streams) and connectionless (for transaction messages).

The Transmission Control Protocol (TCP) became a DOD standard protocol in January 1980. It became a full MIL-Standard in 1983 [35]. TCP provides a reliable, sequenced, flow-controlled channel (virtual circuit) between two processes. TCP is specifically designed to operate above IP and is intended for use in packet-switched networks and interconnected sets of such networks.

Within Air Force and some strategic DOD LAN protocol architectures, it is clear that TCP will be the primary transport layer service; however, some applications may not require all the connection-oriented features of TCP. Other transport services, such as a real-time connection service that offers guarantees on delay and delay variance and a connectionless service that is either reliable or unreliable, may be appropriate for some applications. For example, it has been suggested that some of the functions performed in the Network Voice Protocol would be more appropriate to a general-purpose real-time transport protocol. The User Datagram Protocol (UDP) is a real-time transport protocol that does not impose the

TCP Virtual circuit mechanisms; it simply augments IP with source and destination port addressing and checksumming for error detection.

Internetwork Layer

This performs the end-to-end routing, switching, fragmentation, reassembly and gateway functions needed to interconnect multiple subnets together. This employs a connectionless (or datagram) service.

The primary DOD internetting strategy is embodied in the Internet Protocol. IP is a datagram internetting scheme whereby internet datagrams are routed among hosts and gateways using the local subnetwork's internal routing mechanisms. The services offered by IP thus clearly are the model for the "basic" connectionless internet service as described in the DOD network Layer. IP would sit in the "internet sublayer" of the Network Layer and would call on the services of the protocols within the "subnetwork sublayer" for routing within a subnetwork.

The Internet Protocol (IP), which became a DOD standard in January 1980 and became a MIL-Standard in 1983 [36], uses an extended finite state machine model. The primary services provided by IP are internet routing and packet fragmentation and reassembly. IP provides several features, such as source routing, the "don't fragment" option, and certain "type of service" options, that are not provided by other "potential" network layer protocols such as the NBS Internet Protocol.

DOD must also allow real-time services to be incorporated within its protocol architecture. At the Network Layer, this requires an alternative to IP which can make some guarantee as to the delays and delay variances experienced by user data as it traverses the internet. Such a service can be accomplished via a connection-oriented internet protocol which dedicates gateway or subnetwork resources on a per-connection basis. Although a general-purpose real-time internetting protocol has not been developed within the DOD community, the ST protocol used in the Experimental Integrated Switch Network (EISN) project for voice is a good initial model. Such a real-time internetting protocol would be similar to IP in that it uses the services of the subnetwork-sublayer intranet routing mechanisms for subnetwork transport.

Network Layer - Provides the use of public data network services as one type of subnet under the internet layer.

A type of internetting service which should be incorporable within the DOD Network Layer is reliable connection-oriented service typified by X.75. The degree of "reliability" offered by such systems is generally not enough to eliminate the need for end-to-end mechanisms with the Transport Layer.

Consequently, such an approach to internetting is not anticipated to be used heavily by DOD applications. However, the ability to interoperate with public networks necessitates that the DOD Network Layer not preclude such a service.

Data Link Layer - Provides for the reliable transfer of data units across a link connecting two nodes. This can be multipoint, point-to-point, or a local area network.

Physical Layer - Provides the means to access the physical medium through which data communication occurs.

5.13 ISO and IEEE 802 Networking Reference Models and Protocols

5.13.1 Introduction

The International Standardization Organization (ISO) has established a model of protocol architecture based on a seven layer structure [5]. This model is known as the Open Systems Interconnection Reference Model (OSI/RM).

The basic architecture specified by the OSI Reference Model is only the first in a family of standards that will result from this work. From the established principles of OSI, layer service definitions are now being developed that will then enable further development of the necessary protocols for communications among distributed application processes within the Open Systems Interconnection structure.

The definition of the OSI scope and standards started in the mid-1970's, and a set of standards covering different layers of the OSI model should become available by the mid-1980's. The environment and the objective to be addressed by OSI have been defined as follows [5].

"In the concept of OSI, a system is a set of one or more computers, associated software, peripherals, terminals, human operators, physical processes, information transfer means, etc., that forms an autonomous whole capable of information processing and/or information transfer. OSI is concerned with the exchange of information between open systems (and not with the internal functioning of each individual open system). OSI is concerned not only with the transfer of information between systems, i.e., transmission, but also with their capability to work together to achieve a common (distributed) task. In other words, OSI is concerned with cooperation between systems, which is implied by the expression 'systems interconnection'. The objective of OSI is to define a set of standards to enable open systems cooperation. A system which obeys applicable OSI standards in its cooperation with other systems is termed open system."

OSI defines the external communication capability of a system to make it an open system, i.e., capable of cooperating with other open systems according to OSI standards. The normal OSI applicability is in cases where the operating open systems each have a different internal architecture.

The advent of OSI standards will bring a new dimension to the current environment by providing universally agreed upon means of permitting communication and cooperating between (or among) heterogeneous systems and products. The existing systems will progressively implement OSI capability in response to user application needs. But the existence of OSI standards should not, and will not, slow down the increasing number and diversity of heterogeneous systems and products. In fact, in response to users' requirements, the systems built on heterogeneous architectures will grow in number and in size and they will even provide new functions that are not supported by OSI standards.

The resulting OSI environment will, therefore, be characterized by a large, and possibly increasing, diversity of heterogeneous open systems; heterogeneous because they are built on different architectures; and open because they are capable of cooperating with other systems by implementing the OSI protocols.

The layers of the OSI/RM are described (briefly) in Table 5.13.1.

The upper three layers provide the functions in direct support of the application process, while the lower three layers are concerned with the transmission of the information between the end-systems of the communication. The Transport layer is the essential link between these two groups of functions; it provides end-to-end integrity of the communication, ensuring that the appropriate quality of service from the lower three layers meets the requirements of the upper three layers.

For any system to operate properly, there must be orderly management to ensure all components work in harmony. The management aspects of OSI include the control of initiation, termination, monitoring, and handling of abnormal conditions. There are three management areas that need to be considered.

- Application management - involves: initialization of parameters; initiation, maintenance, and termination of the communication; detection and prevention of OSI resource interference and deadlock; integrity and commitment control; security control; and checkpointing and recovery control.
- Systems management - involves control of the OSI resources and their status across all the layers.

Table 5.13.1. Layers of Open System Interconnection Reference Model

APPLICATION LAYER

Directly serves AP by providing access to the Open Systems Interconnection Environment and provides the distributed information services to support the AP and manage the communication.

PRESENTATION LAYER

Provides the services that allow the AP to interpret the meaning of the information being transferred - syntax selection and conversion.

SESSION LAYER

Supports the dialog between cooperating AP's, binding and unbinding them into a communicating relationship.

TRANSPORT LAYER

Provides end-to-end control and information interchange with the reliability and quality of service that is needed for the AP.

NETWORK LAYER

Provides the switching and routing functions needed to establish, maintain, and terminate switched connections and transfer data between the communicating end-systems (NOTE - The term "network" as used here is a specific OSI technical term and should not be taken as denoting a communications network in the conventional sense.)

DATA LINK LAYER

Provides for the transfer of information over the physical link with the necessary synchronization, error control, and flow control functions.

PHYSICAL LAYER

Provides the functional and procedural characteristics to activate, maintain, and deactivate the physical links that transparently pass the bit stream of the communication.

NOTE: AP denotes applications process.

- Layer management - is concerned with the activation process, error control, and coordination of activities within a layer.

The OSI Reference Model presently deals only with a connection-oriented mode of operation. That is were a path for the communication to follow is established prior to the communication. A connectionless mode of operation also has some popularity and is being prepared as an addendum to the OSI Reference Model. Connectionless communications do not preestablish a path, but route individual units of information as they are sent. The optional datagram service in CCITT Recommendation X.25 is an example of connectionless operation, while virtual circuit service is an example of a connection-oriented communication.

5.13.2 ISO's Open System Interconnection (OSI) Architecture and Protocol Suite

- The Purpose of OSI - To allow any computer anywhere in the world to communicate with any other, as long as both obey OSI protocol standards. The OSI Reference Model is an abstract description of interprocess communications.
- The OSI/RM - Defines types of objects that are used to describe an open system, the general relations among these types of objects, and the general constraints on these types of objects and relations. The document which describes the OSI architecture, ISO 7498, defines these objects, relations, and constraints, and also defines a seven-layer model for interprocess communication constructed from these objects, relations, and constraints.
- A Service Specification - Represents a lower level of abstraction that defines the service provided by each layer. Tighter constraints on the protocol and the implementation that will satisfy the requirements are given. It defines the facilities given to the user and an abstract interface for each protocol layer by specifying the primitives the user would invoke to access the service.
- Protocol Specifications - Represent the lowest level of abstraction in the OSI standards scheme. Each protocol specification defines precisely what control information is to be sent and what procedures are to be used to interpret this control information. The protocol specifications constrain implementations sufficiently to allow open systems to communicate while still allowing differences in implementation.

- Conformance to the OSI Reference Model - Will not assure interoperability among end systems; only conformance to the OSI protocols will assure open systems interoperability.
- OSI Intent - Is not to standardize the internal operation of a system but rather the communication between systems.
- In the OSI/RM - Communication takes place between application processes running in distinct systems in which a system is considered to be one or more autonomous computers and their associated software, peripherals, and users that are capable of information processing and/or transfer.
- In the OSI/RM Architecture - All services provided to the application process users are grouped into seven subsystems of protocol layers, whose entities are distributed among the interconnected systems and which communicate by message exchange among entities of equal layer in the structure according to formal rules of communications called peer protocols. This layering technique is similar to the one used in structured programming, where only functions performed by a module (and not its internal functioning) are known by its users.
- Layer Services - These are the capabilities which a layer provides to a user of that layer. Within a layer, functions will be performed, some of which are not user provided services. Only capabilities seen by the user are the services.
- Layered Functions - These are the activities which are performed within a layer. Functions are performed by layer entities through cooperation with other layer entities distributed throughout the network. Entities of equal layer value communicate by peer protocols.
- Layer Service Access Points - Abstract ports through which users request and receive the services provided by a particular layer.

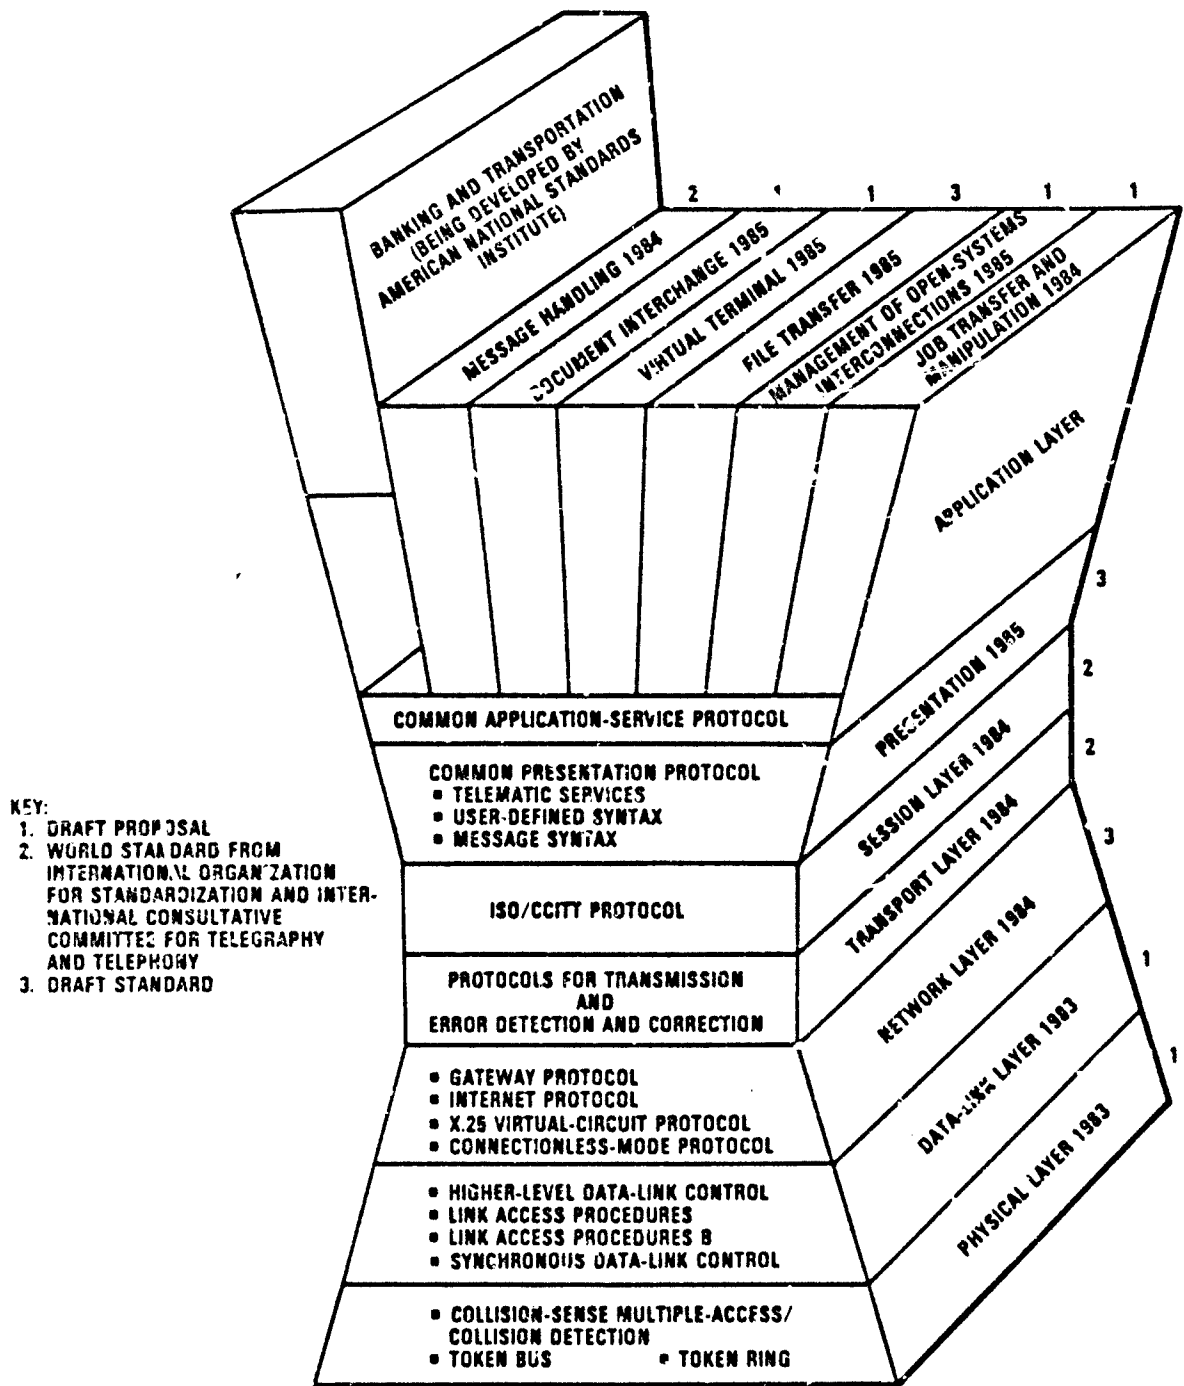
5.13.3 Seven OSI Protocol Layers

This provides a discussion of the seven OSI protocol layers. Table 5.13.3 identifies the corresponding functions of each layer [17]. Figure 5.13.3 illustrates the composite OSI model and status of protocol developments [18].

Table 5.13.3. Functions of the OSI Layers

Functions of the OSI layers	
Application layer	
Common application service elements	
Login	
Password checks	
Set up associations to named peers and agree on the semantics of the information to be exchanged	
Specific application service elements	
File transfer and file access	
E-mail: class virtual terminal	
Forms: class virtual terminal (under development by ECMA)	
Message handling	
Document transfer	
Job transfer and manipulation	
Videotext	
Graphics (semantics)	
Commitment, concurrency, and recovery	
Purchase orders	
Banking protocols	
Credit checking protocols	
Invoice protocols	
Inventory protocols	
	Industry protocols
Presentation layer	
Negotiate transfer syntax for character sets, text strings, data display formats, graphics syntax, file organization, data types, financial information	
Session layer	
Connection establishment and termination	
Data transfer	
Synchronization between end-user tasks	
Graceful and abrupt closure	
Map addresses to names (users retain same name if they move)	
Dialog control (who, when, how long, half or full duplex)	
Quarantining of data (buffering of data until instructed to deliver it)	
Transport layer	
Add s/s end-user machines without concern for route of message or address of machines en route between end-user machines	
Multiplex end-user address onto network	
End-to-end error detection and recovery	
Monitoring of quality of service	
Possibly disassemble and reassemble session messages	
Network layer	
Set up routes for packets to travel (establish a virtual circuit)	
Address network machines on the route through which the packets travel	
May disassemble transport messages into packets and reassemble them at the destination	
Send control messages to peer layers about own status	
Flow control (regulate the rate at which a machine receives messages)	
Recognize message priorities and send messages in proper priority order	
Internetworking	
Data-link control layer	
Add flags to indicate beginning and end of messages	
Add error-checking algorithms	
Make sure data is not mistaken for flags	
Provide access methods for local area networks	
Physical layer	
Handle voltages and electrical pulses	
Handle cables, connectors, and components	
Handle collision detection for CSMA/CD access method	

13457-18



13457-17

Figure 5.13.3. Composite OSI Model

- Application Layer - This deals with the semantics of the applications. It is the uppermost region of the OSI/RM containing not only OSI services, functions and protocols, but the application processes themselves; the users of complete OSI services. The basic OSI services are a set of generic networking-type utilities plus the common application service elements. OSI generic services being standardized are protocols for Virtual File, Virtual Terminal, Job Transfer/Manipulation and Management of Application/System Resources. Other application services beyond these OSI offerings are needed, but industry groups will develop these for their specific needs. In terms of the GNOS object (resource) based architecture model, the OSI generic services provide a set of networking utilities/protocols which support the GNOS resource managers.
- Presentations Layer - This deals with the syntax or data representation for the semantics of the users application data to be transferred through the network. The user may use an existing context or define its own and register it with the ISO.
- Session Layer - This deals with organizing and managing the data being exchanged between application processes, the establishing of synchronization points and the definition of special tokens for structuring exchanges.
- Transport Layer - This provides for the transparent transfer of data between end systems. It optimizes use of any underlying network service and provides the overall end-to-end reliability for communications. Two types of services have been identified; connection-oriented (for data streams) and connectionless (for message transactions).
- Network Layer - This provides for the relaying, routing, switching and internetworking across the heterogeneous subnetworks (among concatenated networks).
- Data Link Layer - This provides the means for transferring data between network entities and detecting and correcting for possible errors. Data links may be multipoint, point-to-point or a local area network.
- Physical Layer - This provides the means to access the physical medium spanning nodes of the OSI system.

- Media Layer - This is not a separate OSI formalized layer but does represent peer protocol functions which are performed by modulated signalling exchanges, to transfer digitized data through a medium.

5.13.4 IEEE Project 802 Protocols Suite for Local Area Networks

This subsection provides a discussion on the IEEE 802 LAN protocols. Reference [41] provided the basis for the findings which follow, in addition to attendance at the November 1983 meetings held in Silver Spring, MD.

5.13.4.1 General

The Institute of Electrical and Electronic Engineers Standards Committee 802 has been working for the last 4 years to develop standards for shared medium local area networks. Four of these standards have been approved and others are nearing completion.

A local network is a system for interconnecting computer, terminal, or peripheral data stations so they can communicate in a local environment - a set of buildings, an office campus, or a manufacturing complex where all of the devices are within a few kilometers of each other. It is possible that several local networks may exist in the same setting.

Any local network will permit a station to attach to a medium for the purpose of transmitting and receiving data. Shared medium local networks are local networks with one further requirement: they must permit several different information processing systems to concurrently use the medium. There is normally no master station or controller of the medium (such as a PBX) in a shared medium local network. Therefore, access to the medium must be autonomous.

The physical interfaces and protocols of shared medium local networks are designed for efficient operation over short distances (a few kilometers) using high-quality media (shielded twisted pair, coax, optical cable) resulting in low error rates (on the order of one in 10^8 bits). Data rates are high - above 1 Mb/s - permitting many data stations (on the order of 200) to share a single local network transmission medium.

The basic motivation for standardizing shared medium local networks stems from the customer's desire to minimize the cost (and duplication) of installing and maintaining several different networks. Shared medium local networks permit different computer systems, each with its own terminals and peripherals, to attach and concurrently use the same physical medium. This is a first step towards the eventual goal of sharing expensive peripheral devices among different computer systems on the same medium. This goal, however, requires standardization of higher layer protocols.

Standards are being developed for baseband and broadband bus media using a contention method called carrier sense multi-access with collision detection (CSMA/CD) (IEEE 802.3), for baseband media using a token ring access method (IEEE 802.5), and for baseband and broadband bus media using a token bus access method (IEEE 802.4). A baseband medium can be defined as a single medium capable of carrying a single information channel. A broadband medium is a single medium capable of carrying multiple information channels, similar to a community access television (CATV) system.

Some standards efforts are further along than others; for example, development has just begun for a metropolitan area network standard (IEEE 802.6). A metropolitan area network is a form of local network that stretches the meaning of "local," since it encompasses a radius of up to 25 kilometers (using CATV or other media).

All of the shared medium and medium access standards have been specified to work under the control of a single Logical Link Control (LLC) standard (IEEE 802.2), capable of providing connectionless and, if required, connection service for any one of the shared media or media access methods.

Finally, a standard for higher layers of shared medium local networks (IEEE 802.1) is being developed to specify a consistent method for internetworking, addressing, and managing local networks and for addressing at (and possibly above) the network layer. This standard will also be used as a companion document to specify the relationship between all of the other IEEE 802 standards.

5.13.4.2 Physical and Link Layers of IEEE 802

ISO has not established standards for LAN physical and link layers. However, for guidance in this area, consideration should be given to the work by NBS, which plans to complete a local area network interface standard for the Federal Government, and to the work by the IEEE 802 Local Networks Standard Committee. Among the international organizations, the European Computer Manufacturers Association Technical Committee 4 is producing local network standards based, in part, on the IEEE 802 work.

The IEEE 802 committee has recently reached full IEEE final approval of several of its Local Network Standards. It defines a multipoint, peer-to-peer communications network, where all communications are a "single hop" without intervening switching/routing nodes. The scope of this LAN standard includes, in OSI terms, Layers 1 and 2 and the medium (see Figure 5.13.4.2-1). Note that the IEEE 802 specification does not define how a LAN IU should be constructed. The specification defines a logical external interface via multiple service access

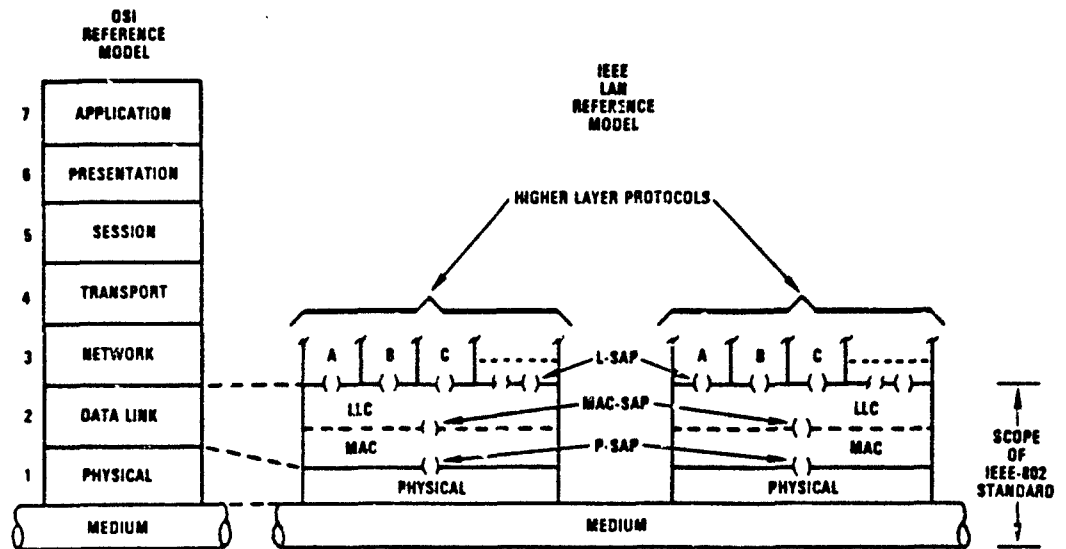


Figure 5.13.4.2-1. IEEE LAN Reference Model

13457-18

points (SAP's) to the next higher (Layer 3) protocol entity and not a physical external interface to a user device.

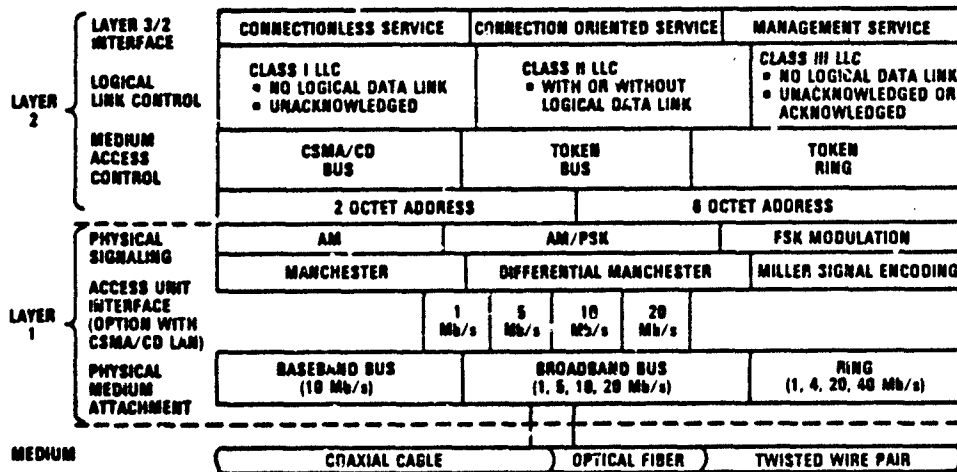
The IEEE 802 Local Network specifications consist of five principal parts:

- Service Access Points (SAP's) - Service Access Points, for addressing end points, are defined. To support multiple Network Layer protocols, LLC Service Access Points (L-SAP's) provide interface ports at the Network/LLC Layer boundary (3/2 interface). The Medium Access Control SAP (MAC-SAP) provides a single interface port at the MAC/LLC boundary. The Physical Service Access Point (P-SAP) provides an interface port to a single MAC entity.
- Logical Link Control (LLC) Sublayer - The Logical Link Control procedures are derived, in part, from ISO's HDLC. Two types of link service are defined: Connectionless and Connection-oriented. Connectionless service allows two LLC stations to exchange frames without establishing a logical link. Frames can (optionally) be acknowledged, but there are no flow control or error recovery procedures. Connection-oriented operation causes a logical link to be established before any exchange of user data. Error recovery, in-sequence delivery, and flow control are provided. A pair of LLC SAP's serve to identify the stream of data units exchanged between

two Network Layer users. Management services (e.g., link layer status reporting) are also defined for this layer.

- **Medium Access Control (MAC) Sublayer** - Two access methods are supported: CSMA/CD and Token Passing. Carrier Sense Multiple Access with Collision Detection (CSMA/CD) is a probabilistic method for assessing the medium. Token passing provides a deterministic method for accessing the medium. The MAC sublayer also provides functions such as frame checking and the discarding of data fragments.
- **Physical Layer** - The Physical Layer provides the capability to transmit and receive modulated signals across the medium. It defines the electrical and mechanical characteristics of the physical medium attachment and physical signaling unit, including the modulation and signaling techniques (e.g., Vestigial Sideband (VSB) modulation, Manchester signal encoding). Each type of medium has a Medium Dependent Interface (MDI). As an option for CSMA/CD networks, the Physical Signaling (PLS) unit may be separated from the Physical Medium Attachment (PMA) by an Access Unit Interface (AUI) cable of up to 50 meters in length.
- **Medium** - Options for the medium include 75 ohm broadband cable, 50 and 75 ohm baseband cable, and 150 ohm twisted wire pairs; fiber optic cable specifications have not been completed.

Figure 5.13.4.2-2 illustrates the options available in the IEEE 802 Local Network Standard.



NOTE: NO RELATIONSHIP SHOULD BE INFERRED BETWEEN OPTIONS AT DIFFERENT LAYERS

Figure 5.13.4.2-2. IEEE 802 LAN Options Available

13457-19

While the IEEE 802 options are extensive, not all the options can be independently selected. For example, selecting one type of medium access method also necessitates a particular medium. Figure 5.13.4.2-3 organizes the physical layer options according to the three primary types of local network technologies (CSMA/CD, token bus, and token ring) recognized by the IEEE 802 LAN standard.

MEDIUM ACCESS CONTROL	CSMA/CD BUS		TOKEN BUS			TOKEN RING	
MEDIUM	BASEBAND COAX 50 Ω	BROADBAND COAX 75 Ω	BASEBAND COAX 75 Ω		BROADBAND COAX 75 Ω	BASEBAND TWISTED WIRE PAIR 150 Ω	BASEBAND COAX 75 Ω
MODULATION	AM	VSB	SINGLE-CHANNEL PHASE CONTINUOUS FSK	SINGLE-CHANNEL PHASE COHERENT FSK	MULTI-LEVEL QUADRINARY AM/PSK	AM	AM
SIGNAL ENCODING	MANCHESTER ENCODING	MILLER ENCODING	DIFFERENTIAL MANCHESTER	"3-SYMBOL" ENCODING	"3-SYMBOL" AND "5-SYMBOL" (OPT) ENCODING	DIFFERENTIAL MANCHESTER	DIFFERENTIAL MANCHESTER
DATA RATES	10 Mb/s	10 (1.5 T80)	1	5, 10	1, 5, 10, 20	1, 4	4, 20, 40

13457-20

Figure 5.13.4.2-3. IEEE 802 - Three Access Methods

5.13.4.3 Baseband, Broadband Standards of IEEE 802

Two forms of CSMA/CD standards are being prepared by the IEEE 802; baseband and broadband.

CSMA/CD Baseband (802.3)

The CSMA/CD baseband network standard specifies an interconnection technique for data stations to share access to a 50-ohm baseband coaxial cable bus. The system corresponds topologically to a branching nonrooted tree. The bus operates at a data rate of 10 Mb/s. Data are impressed upon the bus using a Manchester encoding/decoding technique.

The carrier sense part of the CSMA/CD medium access protocol means that before transmitting a message, a user data station must monitor the bus. If the bus is active, the station must wait. When activity ceases, the station may, after a short delay, transmit its message. The station, however, must also monitor the bus (for a period equal to the propagation time of the bus) to ensure that no other station is also transmitting. This is collision detection.

If no collision is detected, the message is transmitted. If a collision is detected, the data station "jams" the bus by transmitting a

detectable signal, then terminates the message transfer and requeues the message. If collision is again detected, the station increases the delay exponentially, up to a limit. The process continues until a maximum number of collisions is reached, at which point higher layer protocols are advised of the problem.

Each station monitors the bus for its destination address (or an all-parties address or a group address). If a station detects a message with its destination address, it captures and queues the message for a higher layer input process.

Since it is possible to concurrently operate multiple logical data links on the medium, it is necessary to provide both a destination and a source address in the medium access layer protocol. The standard permits two address lengths: a 16-bit address for purely local addressing, or a 48-bit address when (internetwork) addresses of global significance are desired.

The 802.3 baseband CSMA/CD standard was approved by the IEEE Standard Board in June 1983.

CSMA/CD Broadband (802.3)

Efforts to develop this standard are still in their early stages. The CSMA/CD Working Group has begun to specify an interconnection technique for data stations to attach to a broadband coaxial cable bus and to share access to one particular subchannel of the broadband bus.

In order to achieve two-way transmission on a broadband system, either dual cables or a remodulator/translator device is required. A remodulator receives data from any one of the data stations transmitting on subchannel frequency, shifts it to a higher frequency, and retransmits the data down the same coaxial cable. The data stations transmit on the lower frequency. The remodulator is normally located at the transmitter end, or head end, of the coaxial cable. The topology corresponds to the rooted tree (because of the remodulator) with branching.

Two broadband CSMA/CD systems are under consideration. The first would operate at a data rate of 10 Mb/s and is intended to use a larger portion of the existing broadband standard. A bandwidth at least equivalent to two TV channels will be required to transmit the signal over a CATV cable system.

The second system is optimized for broadband operation at a data rate of 5 Mb/s, so that it would require at most the bandwidth of a single TV channel.

Submission of broadband CSMA/CD draft standard to the IEEE Technical Committee on Computer Communications is not expected until 1984 at the earliest.

The Baseband Token Ring Standard (802.5)

The draft token ring standard being developed by IEEE 802 specifies a point-to-point ring topology and a token-passing access method. The token ring baseband standard specifies an interconnection technique for stations to share access on a topological ring. Lower-speed (1 Mb/s to 4 Mb/s) stations are interconnected in a point-to-point manner using 150-ohm shielded twisted pair media, while higher-speed stations require interconnection with coaxial cable or optical fiber. The modulation technique used is Differential Manchester.

The token ring requires that when an operating data station is not originating data transmission, it must repeat the data that it receives. Medium access is controlled by means of a token that is passed from station to station, and grants the holder the right to transmit information on the ring. A station passes the token to its physical successor when it has no more data to transmit, or when a token-holding timer has expired. The token-holding timer prevents one station from hogging the ring.

As each data unit is passed around the ring, a data station that has a message queued for transmission is permitted to make a priority reservation by modifying three bits in a header that eventually indicates the highest priority of messages queued by all other members of the ring. If the message priority of the repeating device is greater than the priority held in the received "reservation" bits, the device may modify the reservation bits to indicate its higher priority request. If the priority of the device is less or equal to the received reservation bits, no modification is made. When the message is returned to the transmitting station (the token holder), the priority of the reservation is noted. When the token-holding station has complete its transmission, either because it has exhausted its priority data or because the token-holding timer has expired, a free token is issued whose priority is set to either the highest reservation received or the priority of the originally received free token, whichever is greater.

Both 16-bit and 48-bit addressing are permitted by the token ring local network standard.

When completed, the IEEE 802.5 token ring draft standard will be sent to the IEEE TCCC for review and ballot. If accepted, the draft standard will be sent to the IEEE Standard Board, which may approve the standard before mid-1984.

Token Bus Standards (802.4)

The draft IEEE 802 token bus standard defines an interconnection technique for devices to share access on a physical topological bus. The standard defines protocols used by the physical and medium access control layers, interfaces between those layers, and interfaces to the medium and to higher layers.

The IEEE 802 token bus standard has been written to include both baseband and broadband systems. The intent of producing a single standard for these two media is to permit an easier transition from baseband to broadband local network systems, with little change in the installed medium or termination equipment.

The IEEE 802.4 token bus baseband and broadband draft standards have already received an affirmative ballot from the IEEE Standard Board.

Baseband Token Bus

This standard specifies an interconnection technique for devices to attach to a 75-ohm baseband truck coaxial cable bus, and to share access to that bus using a token-passing medium access protocol. Bus operations at data rates of 1 Mb/s, 5 Mb/s, 10 Mb/s, and 20 Mb/s are specified. The standard specifies two different modulation techniques. Lower-speed (1 Mb/s) systems use Differential Manchester encoding with phase modulation frequency shift keying. Higher-speed (5 Mb/s and 10 Mb/s) systems directly encode the data using a phase coherent modulation technique.

Either 16 or 48 bit source and destination addresses may be used as station addresses in the baseband token bus medium access protocol.

The token bus medium access protocol is similar to, but must differ from, that used in the token ring since the medium does not form a physical ring. Therefore, the token cannot simply be passed to the next physical data station. The token must be logically addressed to a particular data station, known as the transmitting station's "successor." The transmitting station must maintain the addresses of its predecessor and successor station so it can maintain token-passing operation on the bus in case of failure.

A multiple-level priority mechanism is built into the token bus medium access protocol. But unlike the token ring, where a new priority request can be made as each data unit passes around the ring, the token bus priority mechanism requires that priority be based upon a higher (above medium access) level agreement among the token bus stations, and requires additional individual data unit transmissions among the stations to set up and maintain that agreement.

Broadband Token Bus

The token bus broadband standard specifies an interconnection technique for devices to attach to 75-ohm broadband coaxial cable bus. The data stations use a token access protocol to share a particular broadband subchannel of the medium. Use of two separate physical broadband access cables, one for the data stations originated transmissions, one for head end originated transmission, is also permitted, but not recommended, in the IEEE 802 broadband token bus standard.

The token bus broadband subchannel may operate at data rates of 1 Mb/s, 5 Mb/s, or 10 Mb/s. At 1 Mb/s, one fourth of a standard 6-MHz CATV channel is required to carry the data from the data station to the head end to the data station.

At 5 Mb/s, one standard 6-megahertz tv channel is required to carry the data from the data station to the head end, and another 6-MHz channel is required to carry the data in the other direction. At 10 Mb/s the channel width requirement is doubled.

A single method of data modulation is used in this broadband token system. It is a variant of Duobinary AM/PSK, and requires the encoding and detection of three levels of amplitude that are used to distinguish between symbol codes.

The medium access control protocol used for broadband token bus is identical to that used for the baseband system, with some additional functions required to interface with the head-end remodulator.

5.13.4.4 Metropolitan Area Networks (802.6)

The IEEE 802 executive committee has received permission to expand its charter to write standards for metropolitan area networks operating over distances of 5 to 50 kilometers at data rates at or above 1 Mb/s. Several proposals for MAN standards have been made to the MAN Standard Committee, including systems using broadband cable TV, fiber optics, and packet radio. Possible services to be provided are bulk data transfer, digitized voice, compressed video, videotex, and transaction service. This standard is in the initial definition stage.

It is anticipated that metropolitan area networks may provide access to local networks and also serve as a means of access to satellite or other wide area networks. The approval cycle for the metropolitan area network standard is not expected to begin before 1985.

5.13.4.5 Logical Link Control Standard (802.2)

The standards discussed previously define a physical means to attach a data station to a medium and an access method protocol for sharing the use of that

medium. Any link protocol could be used to transmit data unit messages across the medium, provided it is enveloped in one of the medium access protocols. Because the standard calls for an autonomous medium access protocol, it is possible to operate several independent (logical) links concurrently over the same medium, using the same medium access protocol. The user, however, may want a single data station to be able to concurrently operate on several different logical links through the same single connection to the medium.

If this is the case, the station must have a method to multiplex, demultiplex, and otherwise sort out the data from the multiple concurrent data links that are intended for different users in that station.

The IEEE 802 logical link control standard specifies protocols to control one or more logical links on a single medium, through a single physical attachment of each station to the medium, using a common medium access method. The logical link control protocol uses the following standard protocols: CSMA/CD, token ring, token bus, or metropolitan area network.

The logical link control protocol permits the multiplexing to and from up to 128 distinct logical links in the destination (and source) data station. The number of logical links actually maintained is a function of the resources (buffering and control) available in the data station.

Two forms of logical link protocol service (control) are defined in the standard. Connectionless service is required; connection service is optional.

Connectionless logical link service is similar to datagram service, where the receipt of a link data unit transmission is not acknowledged via the logical link protocol. It is assumed that data units are acknowledged at some higher protocol level, and that retransmissions, when required, are requested by a higher-level protocol. It is assumed that the medium bandwidth is adequate to streamline the transmission of data over highly reliable local networks.

Connection service is similar, in fact, almost identical, to the type of service provided by an X.25 balanced-mode link layer protocol. Acknowledgment of data units and flow control both exist at the link level. Connection service requires higher operational overhead at the link level, but assures that the logical link can operate without overloading the limited buffering capacity of the data station.

Several implementors intend to envelop other (older) link protocols (SDLC, bisync) within the IEEE 802 logical link control protocol in order to provide migration paths to permit older equipment to be multiplexed onto a shared medium local area network.

The IEEE 802.2 logical link control draft standard has been approved by the IEEE Board. The U.S. National Bureau of Standards plans to issue a standard that specifies only the connectionless form of IEEE 802.2 logical link control.

An ad hoc group of implementors met together periodically at the National Bureau of Standards during the past year and decided to attempt to demonstrate interoperability of a common medium. This was accomplished at the National Computer Conference in July, 1984. All systems used an IEEE 802 medium and medium access protocol working under a connectionless IEEE 802 logical link protocol, with CSMA/CD and token bus.

5.13.4.6 Higher Layer Interface (802.1)

When originally established, the function of the higher layer interface group (IEEE 802.1) was to write a companion document for the more detailed IEEE 802 standards to explain the overall intended architecture. But in the process of developing these individual standards, it was determined that a number of similar problems existed in and between these standards, particularly in the areas of addressing, gateways, internetworking, and network management. The work of the IEEE 802.1 group was, therefore, expanded to standardize these areas.

A gateway is a device used to forward data units between two different local networks, or between a local network and a wide area network. The protocols used by the individual (local or wide area) networks may or may not be the same. To transmit data from one of these networks to another, a gateway device must extract data units from the source network, buffer them, and retransmit them onto the destination network. At the same time, differences in protocol must be ironed out at the gateway (which corresponds to the network layer of open system architecture).

The source and destination local network networks may be at the same site, in which case identical addressing structures may be used. Alternately, the two local networks may be at different sites, using different addressing structures that require address transformation. Finally, the local networks may use the services of a wide area network, requiring yet another addressing transformation, as well as a protocol change. These problems are being considered by IEEE 802.1 members as they begin to write the addressing and interconnection portion of their standard.

The architecture portion of the draft IEEE 802 higher layer interface standard (802.1) will soon be sent to the IEEE Technical Committee on Computer Communication for review and ballot. If accepted, this portion of the standard

-- will be sent to the IEEE 802 Standard Board, and approval could occur in 1984.

The addressing and internetworking portion of the IEEE 802.1 draft standard is not expected to enter the review process before late 1984.

5.14 ANSI 100 Mb/s Token Ring LAN Standard

5.14.1 Introduction

This reports on new work under way in the ANSI X3T9.5 Committee developing a standard for a LAN operating ten times faster than the IEEE 802. The standard is called the Fiber Distributed Data Interface (FDDI). This discussion is based upon reference [94].

5.14.2 Discussion of FDDI

The IEEE 802 body now specifies top-end data rates of 10 Mb/s for its local networks. The American National Standards Institute (ANSI), however, is developing network standards for data rates an order of magnitude higher - around 100 Mb/s. These are not futuristic speeds. Most companies involved in developing these standards plan to introduce products supporting such data rates during 1985 and are pushing to finalize the standard by mid-1984. Called the Fiber Distributed Data Interface (FDDI), this proposed new ANSI standard (currently in committee X3T9.5) specifies a token-passing ring architecture for local networks using optical fiber cable.

There are two reasons for the need for high data rate networks: a dramatic increase in computer processing power over the last few years and an enormous increase in the volume of stored or processed data. As a result, it is the lower-speed local networks that quickly become the weak link between devices needing to transfer huge amounts of data as quickly as possible.

Clearly, data rate requirements depend on the services supplied and the network's application. For example, back-end networks, which connect computers to other storage devices and peripherals, require high-speed data transfer. These networks have a fairly small number of nodes (frequently fewer than 50), and span relatively small distances - usually within a computer room. In general, a backbone network has to be at least as fast as the devices on it in order to minimize buffering constraints. As the speeds of hard disks and optical disks increase beyond 40 or 50 Mb/s, back-end networks need to be even faster. In addition, protocols for such networks must provide for "streaming" operations, in which several data packets are sent end-to-end in a single network access. This is essential.

Unlike back-ends, front-end networks typically connect computers to devices that are more user-interactive, such as terminals, word processors, workstations, and printers. Front-end networks can have hundreds of nodes spanning a few kilometers. The IEEE 802.3, 802.4, and 802.5 standards are essentially designed for front-end applications.

Until now users have been satisfied with data rates of 10 Mb/s or so for front-end networks. But if it were available - and affordable - most network users would welcome higher data rates and their services.

For example, a 10-Mb/s data rate is not sufficient to support many real-time voice conversations, much less video traffic. But as the need for these features increases with developments such as teleconferencing, higher data rates seem much more attractive. On a more practical note, low-speed networks are usually adequate for terminals and printers, but for engineering workstations, which require many huge file transfers daily, a high-speed front-end network is almost essential.

The FDDI is a 100-Mb/s token-passing physical ring using fiber-optic cable. The ANSI X3T9.5 committee specifies this as a local network standard. And various corporate architects of FDDI plan to use it for their own products because the standard includes features that support the many applications required for high marketability.

For example, some would use FDDI for back-end computer room applications where the network need only span several meters. Others intend to use FDDI-based products for connections with circumferences of over 100 kilometers. The FDDI specification places no lower bounds on the number of nodes and the distance between them; at the same time, the upper limits are reasonably large, permitting a wide range of implementations. Some of the main limits are:

- Up to 1,000 nodes on the ring
- Up to 2 kilometers between two nodes
- Up to 200 kilometers ring circumference

With the maximum 1,000-node configuration, the average node separation will be 200 meters in order to limit the ring circumference to 200 kilometers. Yet several nodes may be separated by up to 2 kilometers as long as the average separation is 200 meters.

These limits are imposed to minimize latency, or the time it takes a signal to travel around the ring. The maximum ring latency is an important parameter for some real-time network applications. And with the proposed FDDI standard, it is held to only a few milliseconds.

The FDDI ring is a combination of two independent counter-rotating rings, each running at 100 Mb/s. If both rings operate simultaneously, the effective throughput is 200 Mb/s. An FDDI scheme can actually use a special case of this where one ring connects all the nodes, and the second counter-rotating ring only a selected few.

Figure 5.14.2 illustrates a possible FDDI network configuration with fiber-optic cables forming the inner and outer rings. The paths through which the data travels around the ring are also shown. The ring that reaches all of the nodes is called the secondary ring and carries data in the opposite direction of the primary ring. The primary ring connects only the Class A stations (an explanation of station classes can be found below). Such a concentric scheme is useful during ring reconfiguration. If the outer ring fails, for instance, the network can continue operating on the inner ring and still keep the intact portions of the outer ring.

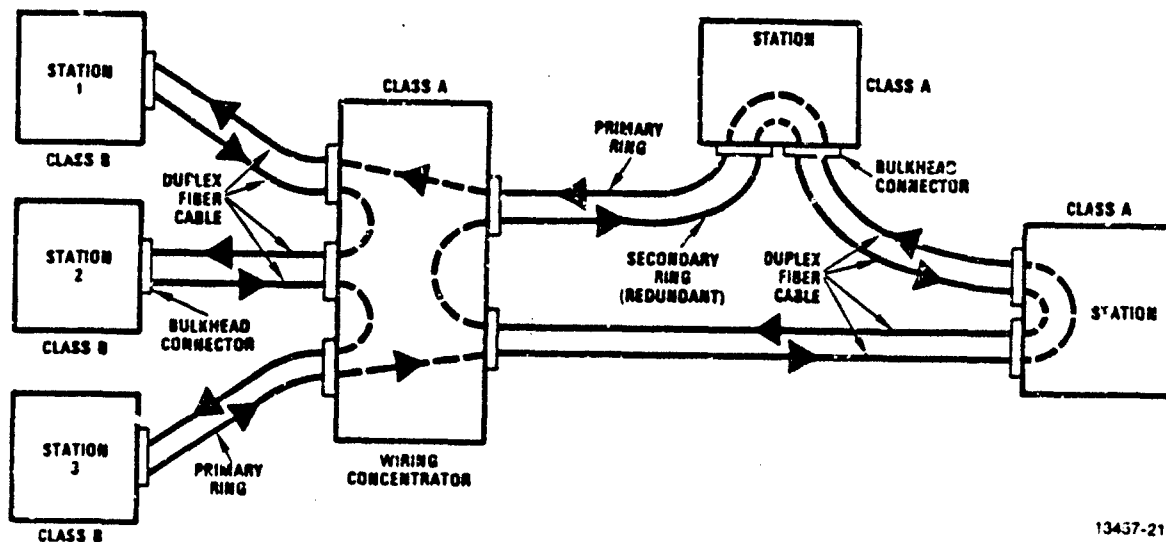


Figure 5.14.2. FDDI 100 Mb/s LAN Topology

Primary and secondary fiber rings transmit lightwave data in opposite directions. Class A stations - either mainframes or wiring concentrators - connect to both the primary and secondary rings. Class B stations connect only to the primary ring. Fiber cable is terminated with bulkhead connectors.

5.16 Air Force Flexible Intraconnect Local Area Network (FILAN)

5.15.1 Introduction

In the last 3 years, Martin Marietta Denver Aerospace has been under RADC contract to design, develop, and document a high-capacity, processor-controlled, bus-oriented local area network (LAN). This effort encompassed the development of a military standard interface specification (MIL-STD-1779) for high-capacity LAN's and a prototype baseline Flexible Intraconnect local area network (FILAN). FILAN is intended to establish the communications foundation for current and evolutionary C³I systems of the future.

FILAN was developed to simultaneously service a wide variety of numerous user equipment, such as processors, peripherals, terminals, displays, voice, video, radar, many military and commercial telecommunications interfaces, wide area networks, and other local area networks. Operational features of FILAN include high availability, flexibility, expandability, and a system manager interface that provides both extensive user services and user-friendly man-machine interface.

5.15.2 System Overview

FILAN is a high-capacity LAN and is shown in Figure 5.15.2. It is designed for use in military C³I systems to provide highly reliable datagram and internet services over a single local subnet (LSN) or interconnection of LSN's dependent on user requirements. Interfaces to the FILAN are in accordance MIL-STD-1779 (USAF). For users not compatible with MIL-STD-1779, adaptation via a programmable interface converter (PIC) is provided.

FILAN is a bus-oriented LAN that employs a prioritized polling scheme to exercise deterministic bus access control. FILAN offers extensive network management capability that allows rapid on-line (re)configuration of the network, with constant performance monitoring, and fault-detection/isolation to protect against component unit (CU) failures and local traffic overloads.

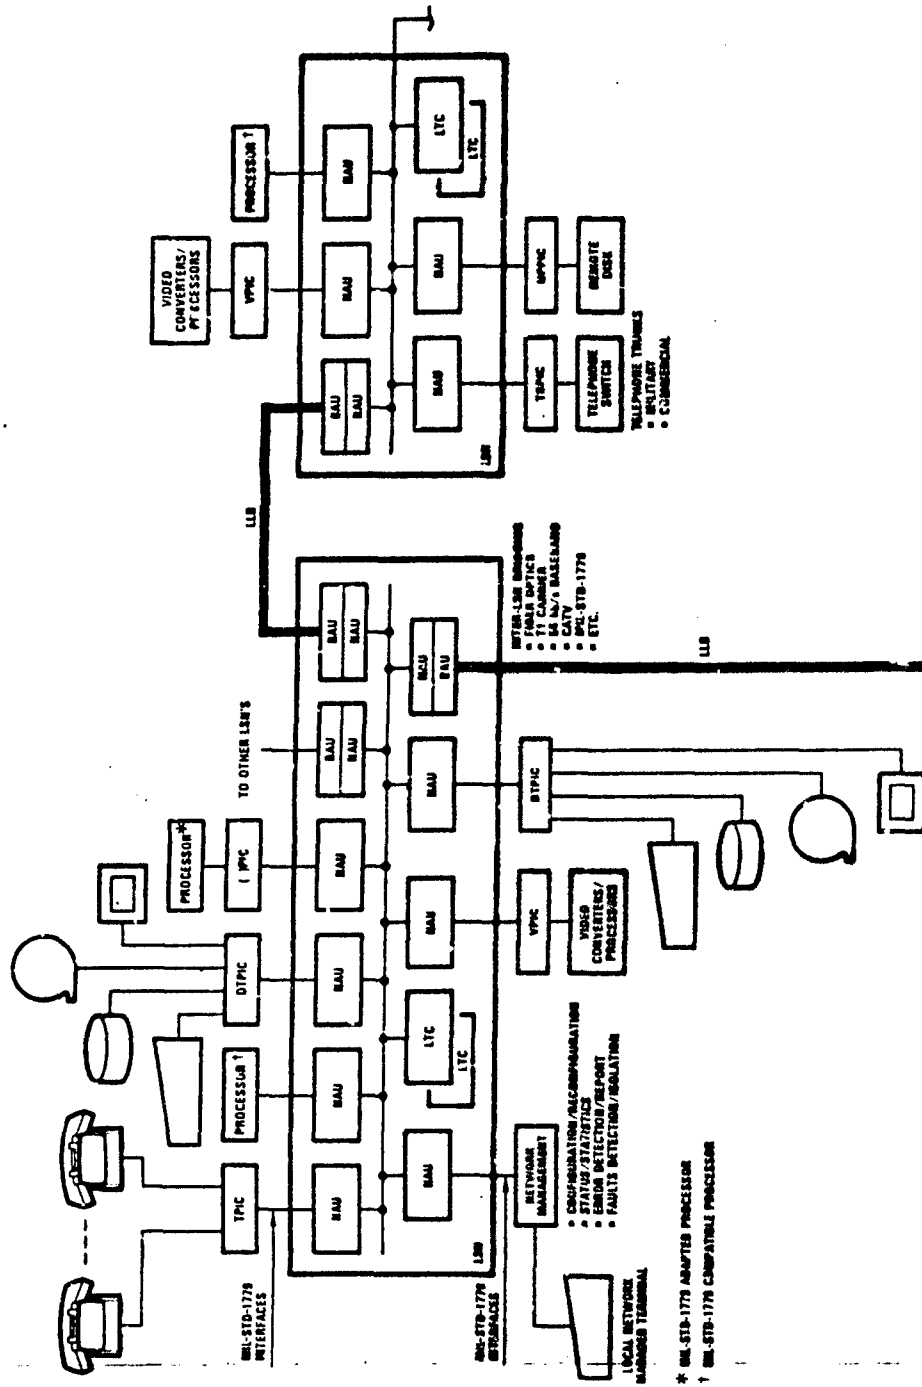
The LSN is comprised of three basic components:

LTC (Local Subnet Token Controller)

NAU (Network Access Unit)

LSM (Local Subnet Medium)

The LTC is the LSN controller - its primary functions are to poll and store the polling schedule and to (re)configure the LSN after the network manager has defined the configuration.



13457-22

Figure 5.15.2. FILAM

The NAU is the connection between the MIL-STD-1779 interface and the bus - its primary functions are to:

- Provide access for multiple virtual users at the MIL-STD-1779 interface
- Provide datagram services to its host users
- Perform MIL-STD-1779 message validation (header contents, header/data error checking and control, message size, and delivery time checking, etc.)

5.15.3 Air Force Workshop on Flexible Intraconnect Local Area Network (FILAN)

This reports on attendance, by invitation, at the Air Force's (RADC) workshop on the FILAN, under development by Martin Marietta - Denver. The workshop was held on March 28-29, 1984 in Denver, Colorado.

A joint Air Force - Martin Marietta conducted workshop was held for the first time to inform people of the current status of the development of the Flexible Intraconnect Local Area Network (FILAN). There were approximately 125 in attendance.

The following were the main results obtained from the workshop:

1. Martin Marietta and Hughes were awarded dual study contracts in 1977 for FILAN. RADC selected Martin's approach and awarded the development contract in 1981.
2. The work to date represents Advanced Development Models with demonstration at 90 Mb/s, with a goal of 180 Mb/s, over ribbon cable and fiber-optic media. These are owned by the Air Force.
3. Extensive hardware and software documentation was developed.
4. A MIL-STD-1779 was developed. This provides a parallel DMA type I/O transfer mechanism for users to access the FILAN.
5. Internally, the FILAN provides a reliable polled datagram access method service encompassing the Physical, Data Link, and Subnetwork layers contained in the Open Systems Interconnection Reference Model. However, the specific protocols employed are not OSI compatible and are not open protocols.
6. Up to 64 Network Access Units (Nodes) are supported, each capable of supporting 14 user ports. Up to 64 FILAN subnets can be combined into a global LAN system.
7. A centralized management with distributed control operation is employed.

8. Bridge links enable interconnecting FILAN subnets together.
9. The current ADM models of the Network Access Units employ 12 Plug-In Circuit Boards in one rack mounted chassis.
10. Users are expected to implement the MIL-STD-1779 interface in the long term. In the interim, Programmable Interface Converters (PIC's) perform the adaptation from the nonstandard to the standard interface. A single PIC requires 13 Plug-In Boards.
11. The current FILAN's NAU's are media independent, but a parallel ribbon cable and a fiber-optic media have been implemented. The Air Force hopes later to standardize a Physical to Media interface(s).
12. Media access is done by a central "deterministic" polled bus access method performed by a Local Subnet Token Controller. The services of point-to-point, multi-point, token ring and broadcast are supported.
13. Traffic handling is stated to be 2000 messages per second per NAU, (1000 in, 1000 out) where each message can be up to 4K x 16 bits of user data (128 Mb/s per NAU).
14. Higher Layer protocols, such as the DOD's TCP/IP, would be treated transparently by FILAN through encapsulation and decapsulation.
15. Network Management currently is done on a VAX 11/780 machine. A user console is provided which enables one to specify and configure the system, reconfigure, and run tests/gather status data. All NAU software is down line loaded.
16. Fiber-optic media work employs a fragmented star, supports 64 stations at 2K meters diameter across the LSN. A laser is employed and uses only an 80 bit turn-on preamble. Expansion to 256 stations was stated.
17. FILAN was stated to be a software intensive system. Forty-thousand lines of code have been developed, about half in FORTRAN and the balance in assembly language.
18. Current FILAN hardware employs 12-13 Plug-In Boards. Power dissipation per NAU is 360 watts. Uses Schottky Bipolar TTL devices. Uses a dual-ported high-speed memory between user I/O and bus media accessing. INTEL 8086, 8089, I/O Processor and Signetics 8 x 300 devices used.

19. Next generation FILAN, called the Basic FILAN Unit (BFU), will employ VLSI, "C" language, IEEE 796 Multibus, Fiber-Optic Media, and new INTEL 186 and 286 devices, aimed at reducing NAU from 12 down to 3-4 Plug-In Boards and might have a cost of about \$10K each. It needs to have MIL-STD-1779 I/F chips developed.
20. There currently is a FILAN test demonstration installed at the NORAD Software Development facility and is undergoing evaluation.

5.15.4 FILAN In a Multi-Media Environment

The FILAN to date is media bound to flat ribbon cable. A fiber-optic link was demonstrated at the FILAN workshop. In a tactical environment other forms of media will be needed to interconnect FILAN Local Subnets and individual nodes together. A new Multi-Media LAN (MMLAN) study has been initiated by RADC with the Harris Corporation to develop a system definition for MMLAN.

5.16 Effects of LAN Protocol Characteristics

This subsection discusses the results obtained from studies of the literature which reported on the effects of LAN characteristics, such as topology, transmission, traffic, access method and throughput-delay performance. The IEEE 802 LAN Medium Access Methods formed the basis of most of the results.

5.16.1 Topological Effects [34]

The common topologies for LAN's are star, ring and bus, plus combinations of these. While there are specific applications for which each of these might be best suited, some general qualities of their topologies for command and control need to be considered.

With a star, the central switching node is a single failure point. The central node is also the limiting factor on throughput and number of virtual connections that can be maintained. A ring with single direction transmission is subject to failure if any node fails because all traffic must pass through all nodes between the source and destination. A ring with bidirectional transmission is subject to fragmentation when more than one node fails. Passive "bypasses" can alleviate the effects of ring node failures, although the "reconstruction" of a damaged ring is difficult because of problems in restoring the token to a known state. Rings allow broadcast transmission and are expandable (but not without temporarily halting operation). Performance is somewhat sensitive to size of the ring.

Bus topologies can have either central or distributed control, with distributed control being more common and generally more useful. Single node failures affect only that node, due to the watchdog timer used to monitor the transmitter. In bus systems using collision detection transmission schemes, intermittent failures tend to be treated as collisions (because both result in checksum error) and are simply corrected through retransmission. Although there are technological limitations, expansion can be achieved with amplifiers, repeaters, and taps up to the bandwidth of the medium. Buses lend themselves either to straight line topologies or tree-like topologies, depending on the transmission technology used and how repeaters or amplifiers and splitters are employed in a particular installation.

Compared to star topologies, buses use considerably less cable and are much more reliable and flexible. Compared to ring topologies, buses are more flexible, offering easier expansion and adaptation to a changing environment. Bus topologies thus have significant advantages for command centers. However, ring networks are capable of supporting high data rates, priority and real/non-real time traffic.

5.16.2 Transmission [34]

There are several media that could be used in LAN's, including twisted wire pairs, coaxial cable, optical fiber, free space (i.e., radio), and infrared. The decision to use a particular medium is usually based on requirements concerning bandwidth, connectivity, geographic scope, noise immunity, security, cost, and suitability for the application at hand. Today, radio is both expensive and not available off-the-shelf; infrared is expensive, not off-the-shelf, and essentially only point-to-point; twisted pair, while inexpensive, is unsuitable with respect to bandwidth (except with multiple twisted wire pairs that may result in space utilization problems), noise immunity, and connectivity.

Optical fiber has the potential to be relatively inexpensive when the technology matures; currently this is expensive and just becoming easily available. Optical fiber has complete noise immunity, does not emanate, and possesses superior bandwidth capability. Tapping optical fiber is difficult - this is a plus (from a security point of view) and a minus (from an installation point of view). At present, only baseband transmission techniques have been fully developed for optical fiber, which restricts its use to a single information channel.

Coaxial cable for LAN's makes use of established cable television technology, which accounts for its cost advantages and easy availability of components. Other characteristics of coaxial cable used in a LAN depend on whether baseband or broadband signalling is used.

With a baseband system, an attached device transmits bidirectionally along a single baseband cable. Broadband systems are directional, with the interface unit broadcasting (on a "reverse" channel) to the "head-end." The head-end retransmits the signal on the "forward" channel to all other devices.

Broadband systems may use either single or dual cables for transmission. A mid-split broadband system allows the interface unit to transmit and receive at different frequencies on the same cable. This is accomplished by a frequency shift at the central retransmitter. Dual cables allow the interface unit to transmit and receive at the same frequency on different trunks. The full bandwidth of the cable, typically 5 to 300 Megahertz (or 400 MHz), is available. Dual cable systems use unidirectional amplifiers, splitters, and taps; bidirectional transmission components are necessary for mid-split systems. Balancing of signal levels throughout the network is somewhat more complex for a mid-split system.

Compared to broadband, baseband signalling currently is capable of somewhat higher data rates, and the interface units are less expensive, but it is limited to a single information channel, has less noise immunity than broadband systems, and can span distances of only 1 to 3 kilometers. Broadband systems support topologies spanning 10 or more kilometers and multiple video, voice, and data channels, up to a total bandwidth of 300 Megahertz (or more, depending on the cable type). Broadband signalling is highly dependent on modem quality; the modems are principally responsible for the current cost differential compared to baseband.

In summary, broadband cable systems appear to offer significant advantages for command center LAN's. These advantages include the capability to support multiple information channels and types and the ability to cover a larger geographic area (than baseband).

5.16.3 Traffic Effects [34]

Throughput Requirements. LAN throughput requirement is a function of the capacity required of the transmission medium plus the processing capacity at each network node. A quantitative analysis of LAN applications is needed to determine:

- Message sizes (peak, average, and variances)
- Responsiveness characteristics
- Interval between messages (peak, average, and variances)

Traffic Classes. It is desirable that the LAN support both connection-oriented (stream) and connectionless (datagram) classes of traffic. Initially, however, most LAN applications will require a connection-oriented transport service such as the Transmission Control Protocol (TCP) provides. For example, a reliable end-to-end transport service is needed for applications such as file transfers to/from a local word processing workstation or bulk data transfers between cooperating processes residing on different host processors. Depending upon the security architecture, transaction-oriented traffic such as data base query and response may require the establishment of connections. Nondigital (voice and video) stream traffic also requires connection-oriented transport services. Connectionless services might be useful for such applications as the LAN accounting functions.

Tolerable Network Delay. Another LAN parameter is the network-induced delay that can be tolerated by the various LAN applications. Requirements can be expressed in terms of absolute delay (maximum allowable delay experienced during transport) and delay variance (the maximum allowable delay variation within a data stream).

In a local area network (as opposed to a long-haul network), most of the network-induced delay results from protocol processing (as opposed to transmission delays). Since there is a tradeoff between reliability (achieved by error detecting and correcting protocols) and delay (the performance penalty extracted by protocol processing), the LAN designer should note that the command center will put a significant processing load on the LAN nodes due to its stringent reliability (and also security) requirements. Inadequate processing power within the LAN nodes will lead to bottlenecks within the LAN. There are tradeoffs in these two approaches between the nodal processing resources used and the transmission bandwidth resources used.

Interconnectivity Between Nodes. Necessary tools for the LAN design development are internodal data flow models for the major applications. The models must include the data rates (average and peak) and other traffic characteristics (such as stream or transaction traffic). These flow models are needed to determine protocol processing capabilities required in each node and the adequacy (e.g., in terms of connectivity and reliability) of various LAN topologies.

Within the security constraints, the C² LAN will provide full (physical) connectivity between all nodes. However, the local system administrator and the security officer need the capability to block (logically) certain connections (for example, to ease the burden on an overloaded host processor or in response to a security violation).

Concurrency of Connections. Another LAN design parameter is the degree of concurrency of network connections that terminate at a given node. The communications hardware and software must have the necessary capacity to meet the traffic processing requirements of the applications supported by the node. For example, the buffer space available within a node must be sufficient for all virtual circuits the node may establish during a period of maximum connection concurrency. Thus, it is important in the LAN design to establish the maximum and average number of concurrent connections at each node.

Information Transmission Summary

The types of traffic - digital data, voice, video, and images - that may be transmitted on the LAN place two major requirements on network services:

- Transmission rates must be adequate for each information type.
- Protocols appropriate to each information type must be provided.

Without a broadband local area network, multiple cable will be required to support the distribution of video, voice, and digital data. As a point of reference, the EIA 40.1 group, in conjunction with the IEEE 802 Local Area Network committee, is attempting to define how frequency channels on a broadband LAN might be allocated. Draft C of the IEEE 802 Local Network Standard defined thirty-four 6 MHz channels, using conventional CATV nomenclature for the channel numbers. Four 1 Mb/s broadband channels would be allocated to one 6 MHz channel; 5 Mb/s broadband channels would use one 6 MHz channel; a 10 Mb/s broadband channel would use two adjacent 6 MHz channels (or possibly one if four level Vestigial Sideband (VSB) modulation is used). A pairing of forward and reverse channels is also suggested for mid-split single trunk configurations.

5.16.4 Performance Effects of LAN Access Methods

The literature was reviewed to obtain results of previous performance comparisons done for LAN media access protocols. Summary results are presented below on several access methods and the IEEE 802 LAN protocols.

5.16.4.1 Contention Access Techniques

Report [34] provided the basis for the following results.

ALOHA:

This was originally developed for satellite packet radio. Messages can be transmitted at any time. A random timer is started when a message is transmitted; if it is not acknowledged before the timer runs out, it is retransmitted.

Throughput - about 18 percent of circuit capacity

Stability - tends to become unstable causing frame delay increased and throughput decreased at loads above 18 percent

Slotted ALOHA:

A modified ALOHA scheme whereby a message can only be transmitted at beginning of a specified time slot.

Throughput - about 36 percent of circuit capacity

Stability - also tends to become unstable at high loads

CSMA and CSMA/CD:

Carrier Sense Multiple Access (CSMA) requires that a node listen to a carrier signal before transmitting. If another node is transmitting, the node will either back off for a specified time interval before listening again or continuously monitor the carrier signal until the network is clear to send. Collision Detection (CD) is a mechanism whereby if a collision occurs during transmission, both sending nodes back off different time intervals (e.g., an exponential back off algorithm) and try again.

Throughput - CSMA is approximately 85 percent for small networks; increasing propagation delay decreases this value relative to the data frame length. CSMA/CD increases maximum utilization to 98 percent of the circuit capacity. Performance decreases with decreasing frame size and increasing propagation delay between the farthest transmitters. For 10 Mb/s or greater networks, the capacity of CSMA/CD drops to 66 percent.

Stability - CSMA does exhibit some instability at high loads. CSMA/CD exhibits stability under extreme overloads.

5.16.4.2 Deterministic Access Techniques

Report [34] provided the basis for the following results.

Token Passing:

A control token scheme where a special message is passed around the network and the node with the token is allowed to transmit.

Throughput - Depending on the nodal processing delay, the capacity of a token passing network approaches 98 percent. Above 10 Mb/s, this capacity is not affected. Large size networks with lots of nodes or small packets may increase nodal delays significantly. Transmission delays may be two to three times greater than CSMA/CD if nodal processing delays are large.

Stability - is very high.

5.16.5 Comparing Three IEEE 802 LAN Systems Performance

A subcommittee of the IEEE 802 has prepared a report [52] on expected performance of three of the types of media access control systems specified by the IEEE 802 draft standards. That report indicates that the CSMA/CD systems can be expected to yield the shortest delay under light loading, but that the token bus and token ring systems give superior performance under moderate to heavy loads (see Subsequent Discussions).

In their analysis, the subcommittee did not consider the built-in priority functions of the token bus and token ring systems, since they were not yet defined. In light of this added capability, the token bus and token ring local network systems must be considered superior to the CSMA/CD systems for those applications where data of different priority must be transmitted across the same local network.

Because medium access and logical link LSI devices are now becoming available for CSMA/CD systems, these will be the dominant form of commercial shared medium local networks during the next few years. LSI devices for token ring and token bus access will be available in the next 1-2 years.

Token bus baseband and broadband systems seem to be the choice of the industrial automation users, based upon the work of the U.S. Process Data Highway Committee (PROWAY) of the International Electrotechnical Commission (IEC) standard body. Subcommittee 65, Working Group 6 of the IEC is presently extending the IEEE 802 token bus draft standard for use in industrial environments.

Token ring baseband systems have been chosen by IBM and several other major computer manufacturers as their preferred LAN because of the token ring multipriority level capability, because the token ring provides a deterministic, rather than statistical, access time, and because as data rates increase, the required dead time between transmissions on a LAN ring is shorter than that required for comparable bus topologies.

5.16.5.1 Detailed Comparisons [52]

A highlight of the IEEE 802 draft report of the subcommittee is a numerical comparison for the different media access methods with a fixed common message transmission load. Two workloads were used to bound performance: in Figures 5.16.5.1-1 through 5.16.5.1-3, one station out of 100 is always actively transmitting a message and message size is varied: 500, 1000 and 2000 bits. In Figures 5.16.5.1-4 through 5.16.5.1-6, 100 out of 100 stations are active with message transmission. The horizontal axis or abscissa as in either case is the raw data transmission speed of the network (the clock rate), while the vertical axis or ordinate is the actual carried data rate, the rate at which data bits are successfully transmitted. The ideal case is to use zero transmission capacity for network access control, which would be a straight line with slope unity. The deviation from this straight line shows the penalty paid using the same network to control access as well as transmit messages.

Figure 5.16.5.1-7 is a plot of lower bound on message delay versus number of active stations, assuming each station goes idle for a mean amount of time T_{idle} and then active in order to transmit a message.

The best available evidence today is:

- Token passing via a ring is the least sensitive to workload and offers short delay under light load and controlled delay under heavy load.
- Token passing via a bus has the greatest delay under light load and under heavy load cannot carry as much traffic as a ring and is quite sensitive to the bus length (through the propagation time for energy to traverse the bus).
- Carrier sense collision detection offers the shortest delay under light load, while it is quite sensitive under heavy load to the workload and is sensitive to the bus length (the shorter the bus the better it performs) and to message length (the longer the packet the better it does).
- Reservation access via a bus is the least understood of all access methods, yet may offer the simplicity of access under light load of carrier sense collision detection, and the controlled access under heavy loading of token passing.

5.16.5.2 Delay VRS Throughput for LANS

This provides comparison of CSMA/CD and token ring delay versus throughput characteristics. Report [53] provided the basis for the results given.

MAXIMUM MEAN CARRIED DATA RATE VS ACTUAL TRANSMISSION RATE

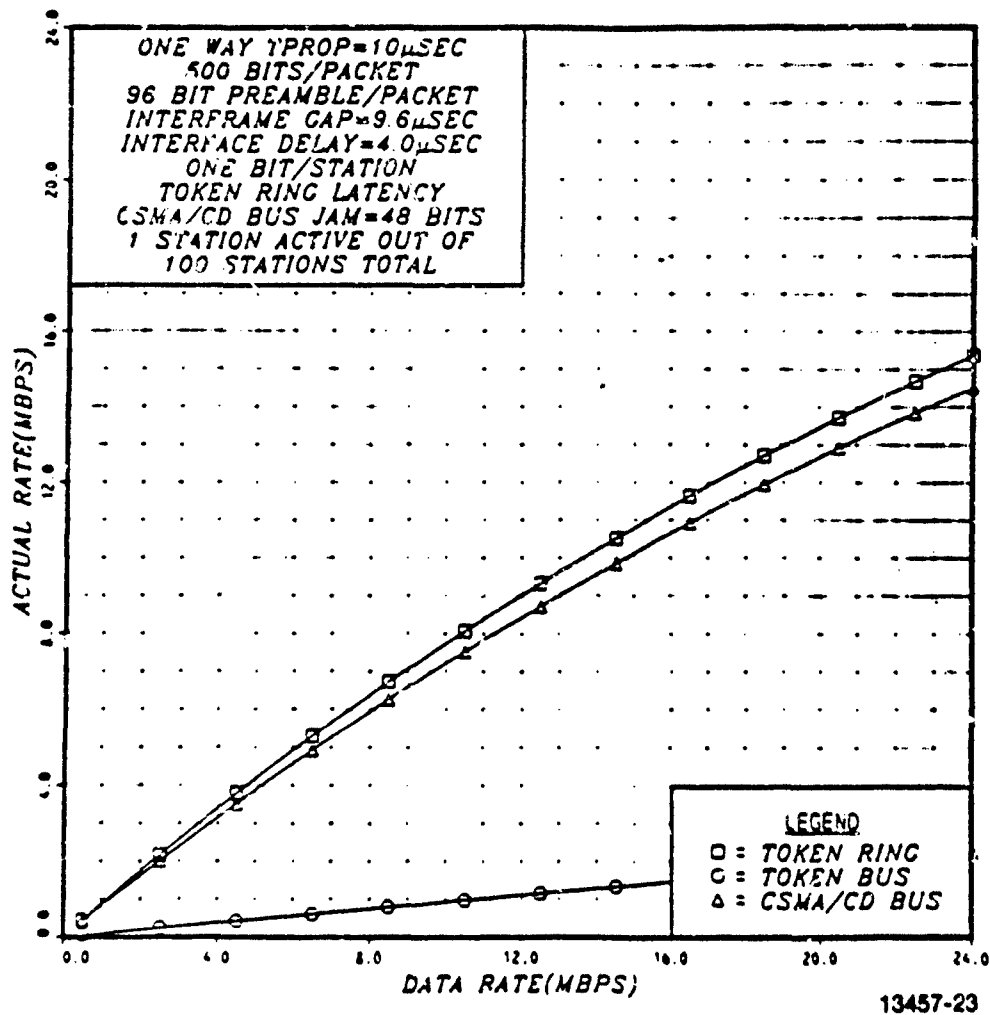


Figure 5.16.5.1-1. Maximum Mean Carried Data Rate Versus Actual Transmission Rate (500 Bit Packet and One Active Station)

MAXIMUM MEAN CARRIED DATA RATE VS ACTUAL TRANSMISSION RATE

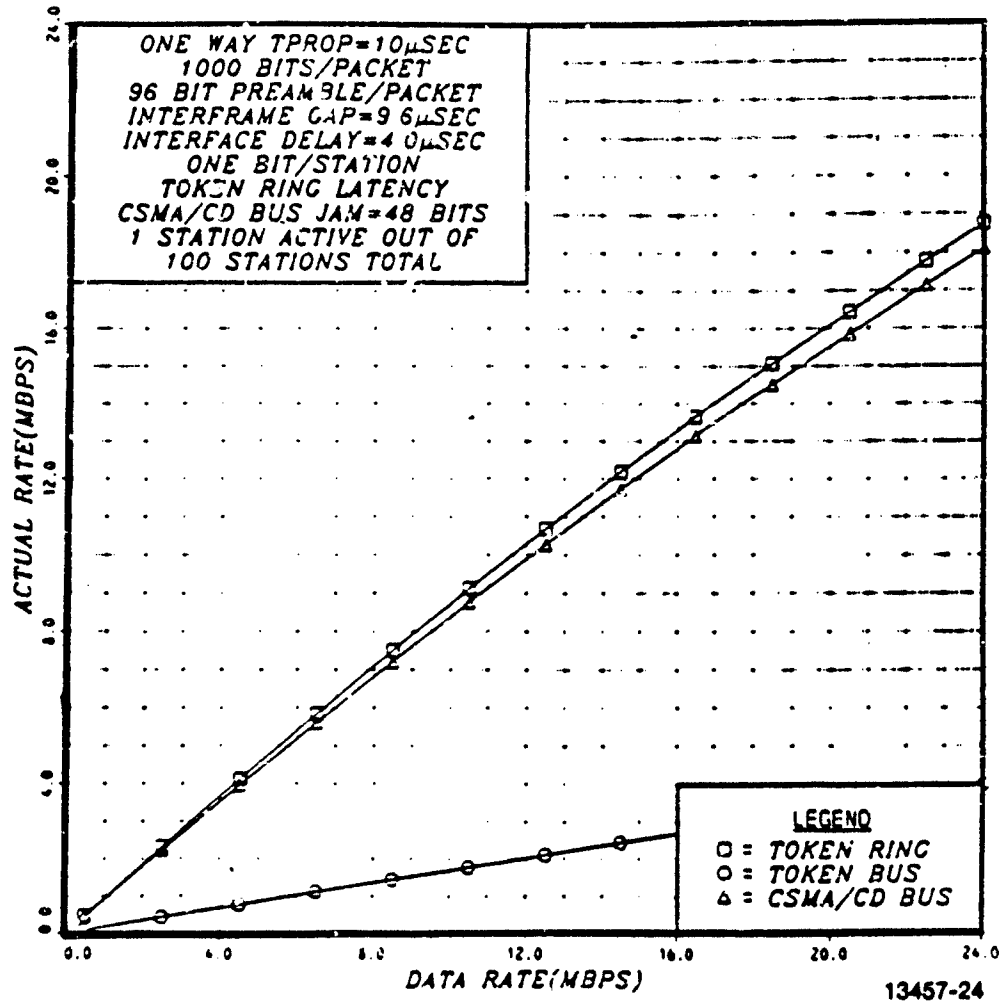


Figure 5.16.5.1-2. Maximum Mean Carried Data Rate Versus Actual Transmission Rate (1000 Bit Packet and One Active Station)

MAXIMUM MEAN CARRIED DATA RATE VS ACTUAL TRANSMISSION RATE

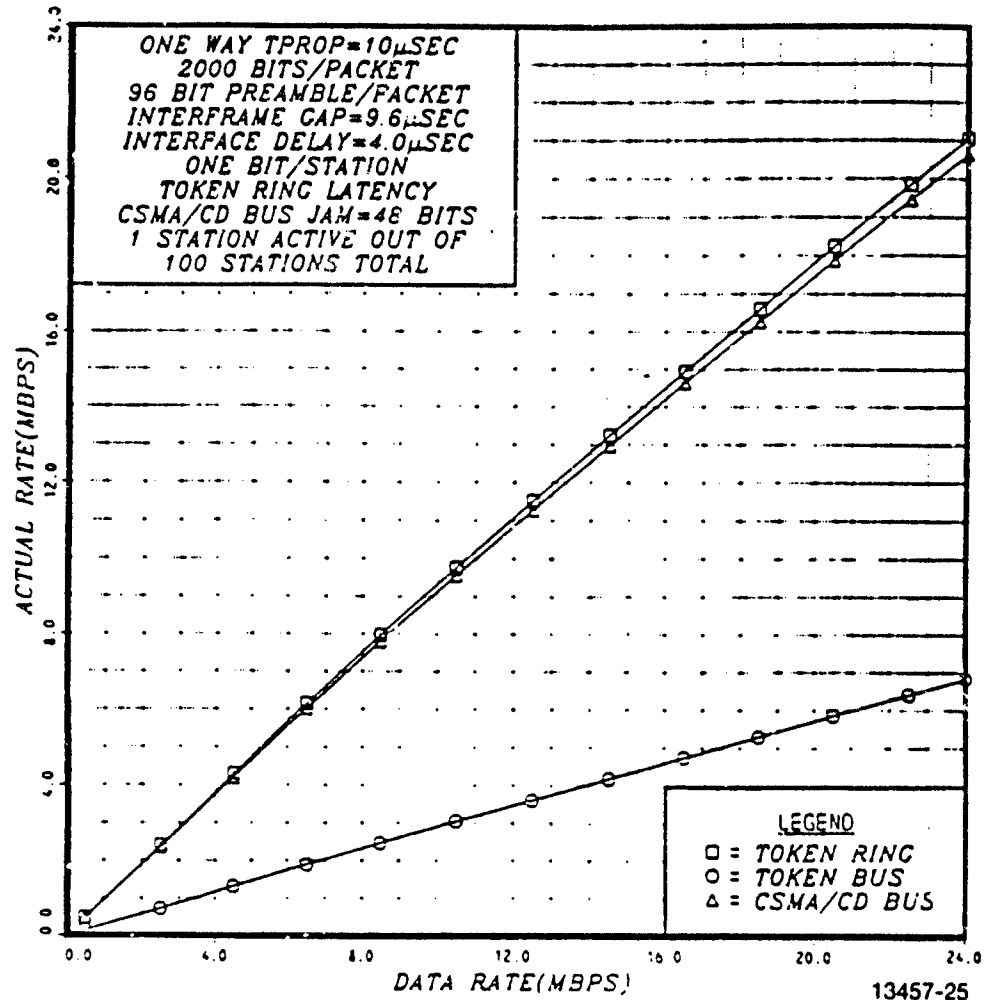


Figure 5.16.5.1-3. Maximum Mean Carried Data Rate Versus Actual Transmission Rate (2000 Bit Packet and One Active Station)

MAXIMUM MEAN CARRIED DATA RATE VS ACTUAL TRANSMISSION RATE

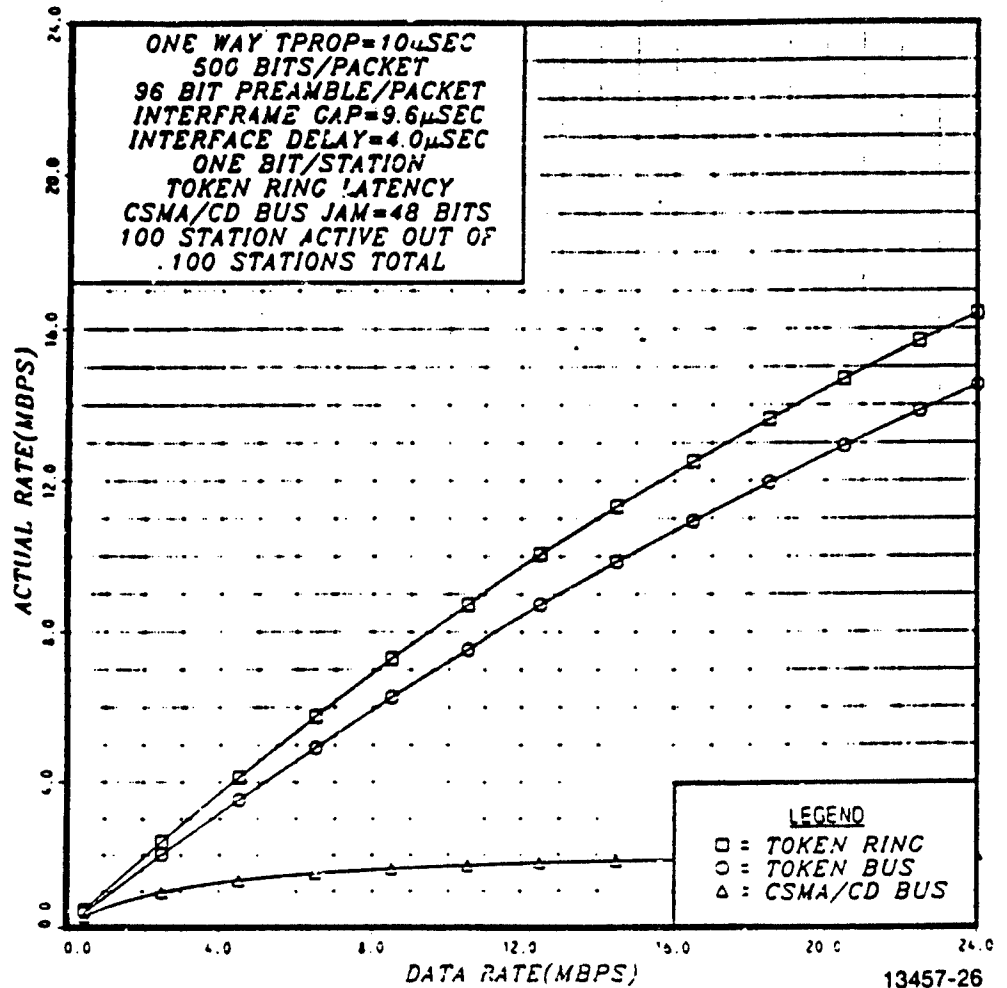


Figure 5.16.5.1-4. Maximum Mean Carried Data Rate Versus Actual Transmission Rate (500 Bit Packet and 100 Active Stations)

MAXIMUM MEAN CARRIED DATA RATE VS ACTUAL TRANSMISSION RATE

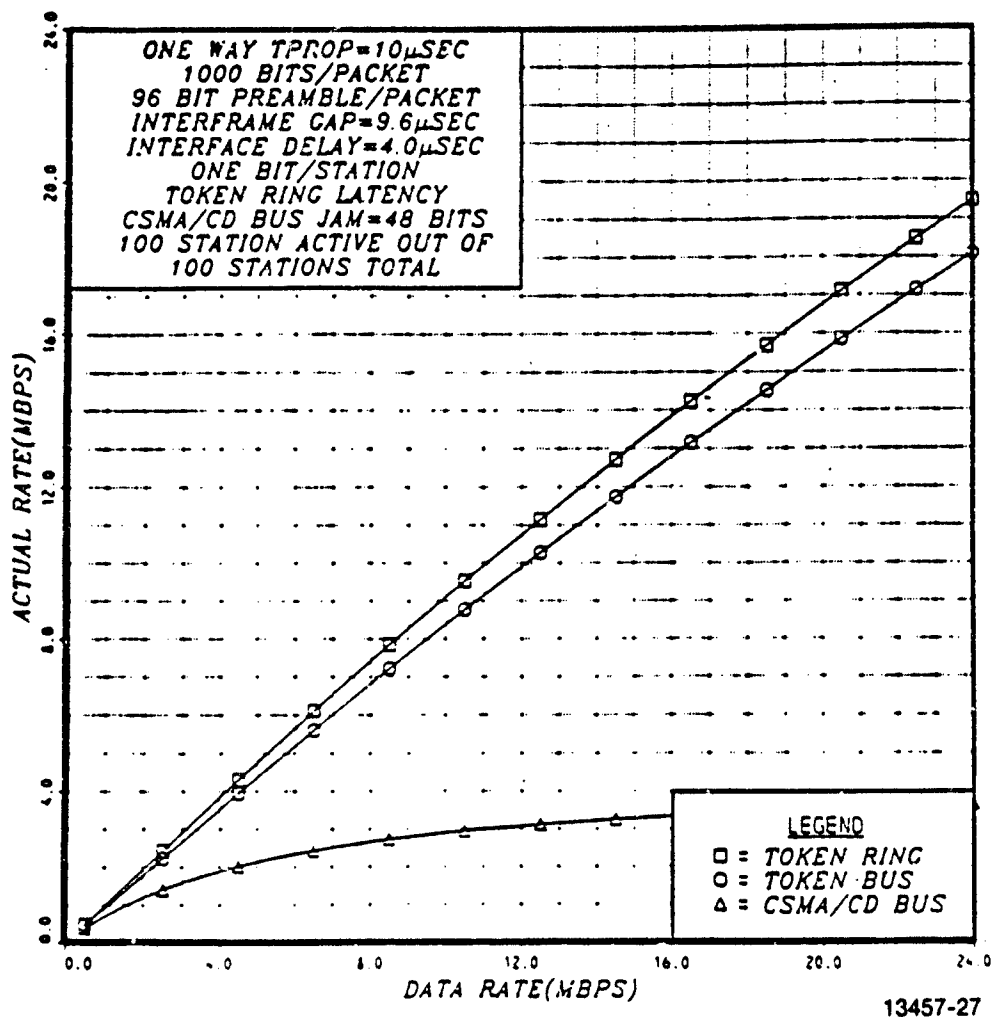
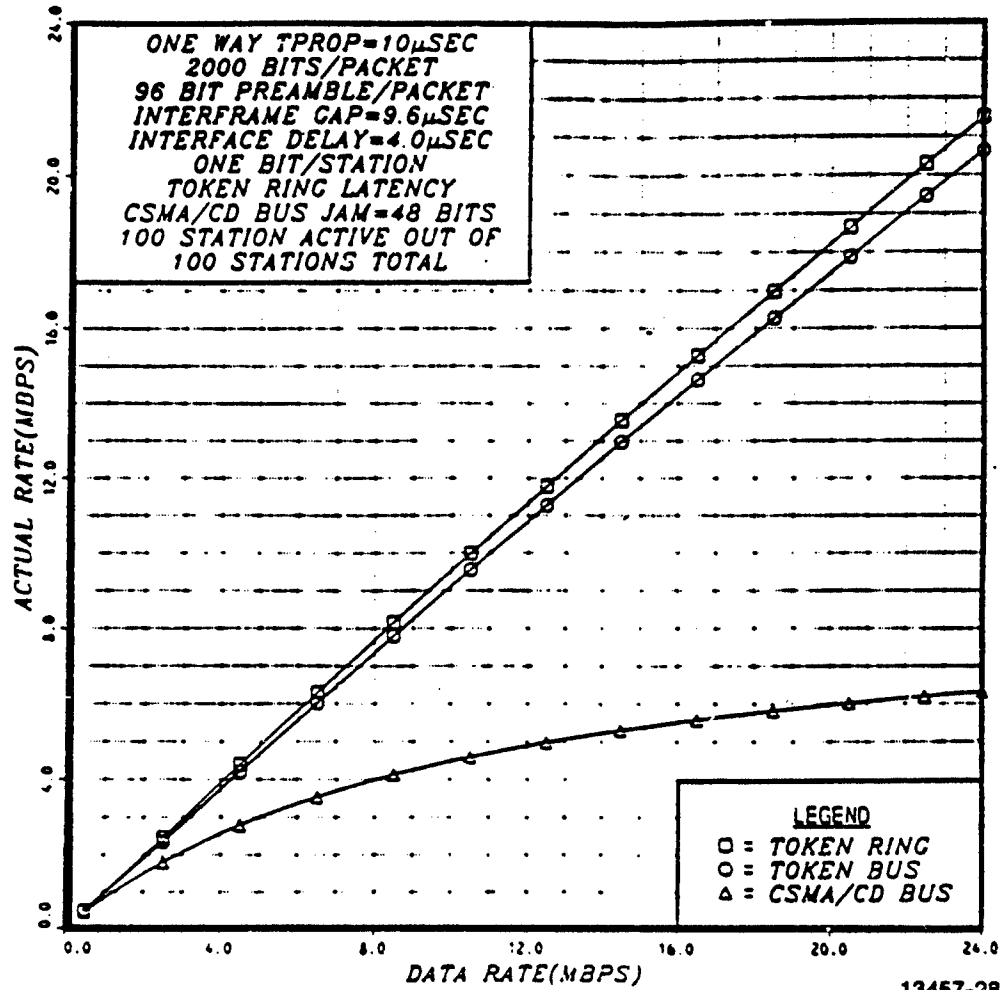


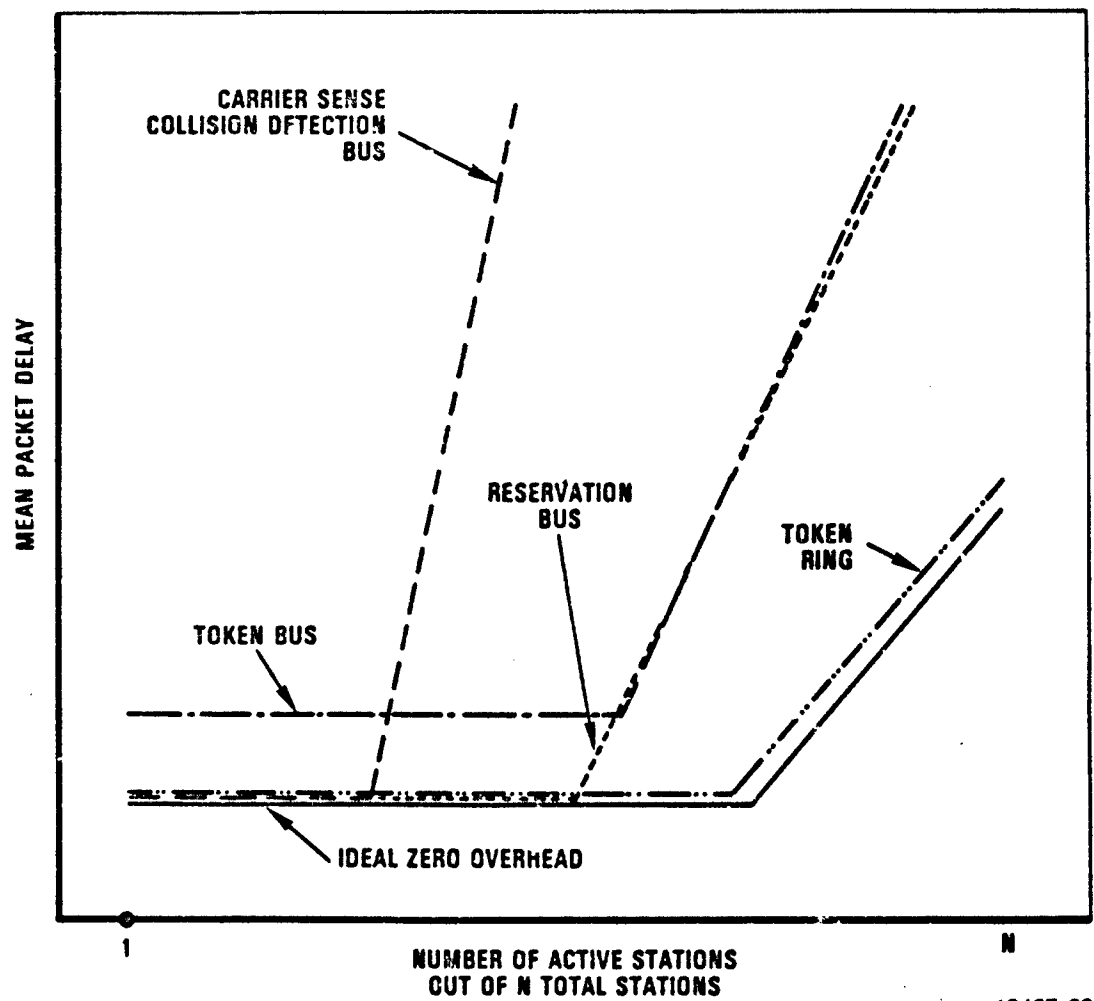
Figure 5.16.5.1-5. Maximum Mean Carried Data Rate Versus Actual Transmission Rate (1000 Bit Packet and 100 Active Stations)

MAXIMUM MEAN CARRIED DATA RATE VS ACTUAL TRANSMISSION RATE



13457-28

Figure 5.16.5.1-6. Maximum Mean Carried Data Rate Versus Actual Transmission Rate (2000 Bit Packet and 100 Active Stations)



13457-29

Figure 5.16.3.1-7. Mean Packet Delay Versus Number of Active Stations

Two performance aspects are of primary interest: the delay-throughput characteristic of the media-access control schemes, and system behavior when the load approaches the saturation point. There exists a large body of performance analyses of ring and bus systems. Some examples are given in [54, 55, 56]. In the following, use is made of some of the results reported in [57]. Figures 5.16.5.2-1 through 5.16.5.2-3 [54, 55, 56] show the delay-throughput relation of token ring and CSMA/CD bus for two data rates: 1 Mb/s and 10 Mb/s.

The general conclusions we can draw from these results are: 1) at a data rate of 1 Mb/s, both systems perform equally well; 2) if the data rate is increased to 10 Mb/s, the token ring has better performance characteristics over a wide range of parameters. In Figure 5.16.5.2-1, the frame-length distribution is negative exponential with an average value of 1,000 bits. A frame represents the entity transmitted by a station when it has access to the medium. The critical parameter that determines the performance of the CSMA/CD bus is the ratio of propagation delay to mean frame transmission time. Since the propagation delay is independent of the data rate, this ratio increases with the data rate. Theory shows [56] that a CSMA/CD bus behaves ideally as long as this ratio is sufficiently low. If, for reasonable traffic loads, it exceeds 2-5 percent, the increasing collision frequency will cause significant performance degradation.

If on a CSMA/CD bus, collision occurs, transmission will be aborted, and the station will reschedule its frame by selecting a random retransmission interval, the length of which is dynamically adjusted to the actual traffic load to avoid an accumulation of retransmissions. The high collision frequency at high load levels together with the retransmission policy causes the variation of the transfer delay to grow. The practical consequence is the danger of stations becoming locked-out for an unpredictable period of time. A token ring, on the other hand, guarantees fair bandwidth sharing among all active stations even at high load levels because the token has to be relinquished after the transmission of one frame.

The general validity of the conclusions drawn above is supported by Figures 5.16.5.2-2 and 5.16.5.2-3. In Figure 5.16.5.2-2, all parameters are the same as before except for the length of the cable, which is now 10 km instead of 2 km. The curve for the CSMA/CD bus at 10 Mb/s illustrates the impact of the propagation delay, and confirms the importance of the ratio propagation delay and average frame transmission time. As a practical consequence, all CSMA/CD systems being discussed specify a maximum distance which is less than 10 km. Finally, Figure 5.16.5.2-3 further demonstrates the robustness of the results. There, the frame-length distribution has a coefficient of variation of 2.

Delay-Throughput Characteristics for Token Ring and CSMA/CD Bus

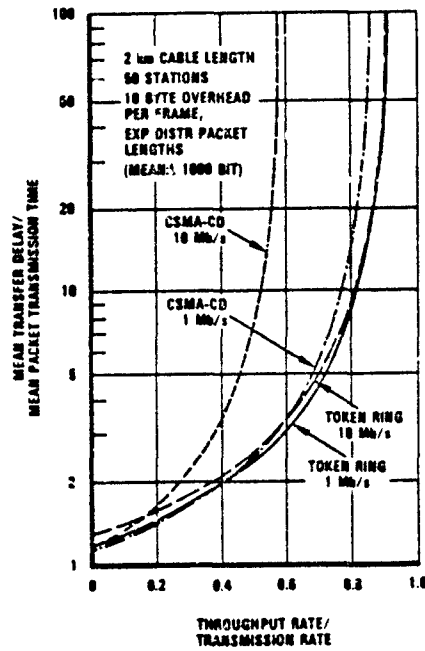


Figure 5.16.5.2-1.

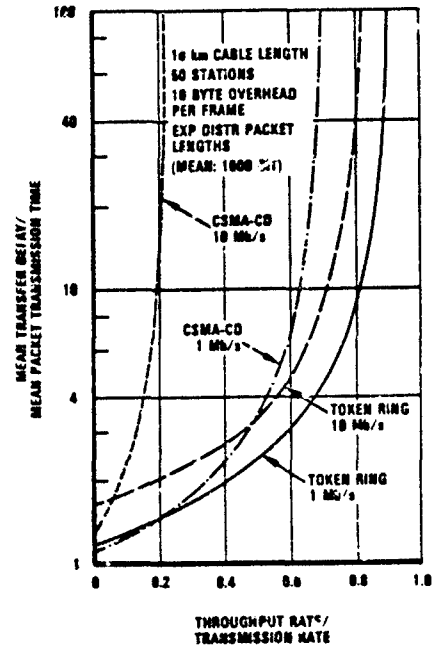


Figure 5.16.5.2-2.

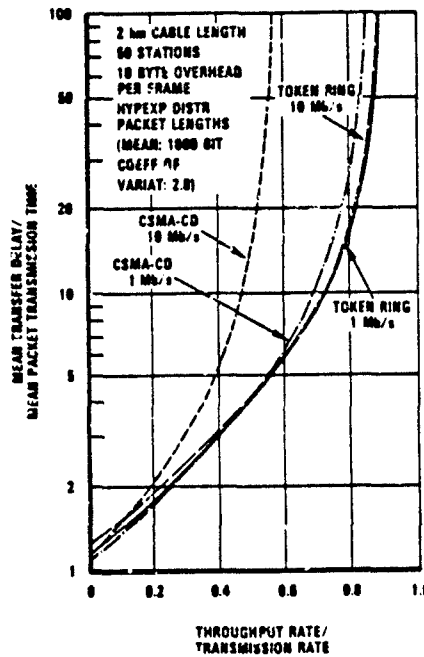


Figure 5.16.5.2-3.

13457-30

5.17 Evaluation of TCP/IP in a Local Network Environment

5.17.1 Performance Results

Review was made of experiments previously conducted by MITRE on TCP/IP in a LAN environment. Reports [45, 46, 47, 48] were used as the basis for the findings reported here.

It was considered essential in reference [45] for the program-to-program communication performance of a LAN-based command center to be comparable to the process-to-process speeds found in existing mainframe hosts. Evaluations of a prototype network, using a media signalling rate of 890 kb/s was made [45, 46]. (The 890 kb/s rate was limited by the peripheral interface chips used.) Employing 512 byte data messages resulted in a 348.9 kb/s rate between TCP users. It was noted in the report "that this rate is almost 10 times greater than any previous TCP implementation." Fine tuning of TCP retransmission and acknowledgment policies were modified more than any other to achieve these results.

Subsequent to the above reporting on TCP performance, reference [47] reported another set of results. This involved using a 10 Mb/s Ethernet CSMA/CD configuration on a 68000 board containing TCP/IP. At a 10 Mb/s media signalling rate, and a 4 millisecond packet processing time limiting factor (caused by the processor on the Ethernet board), an effective transfer rate from source to sink of 892 kb/s was achieved. It was concluded that performance of TCP/IP was good, the number of bytes of overhead was not an important issue, that TCP/IP had built-in security features and that 1 Mb/s throughput would be achievable with the right hardware.

Another MITRE investigation [45, 48] considered the possibilities of subsetting TCP/IP when used for intra-LAN traffic exchanges. A "discretionary" set of TCP/IP transport protocols was considered. It was reported that "they provide a flexible means of using IP and TCP for high bandwidths, low delay environments and at the same time preserve the ability to gracefully, via a local long haul gateway translator, interact with computers located on a long haul network."

This approach was based upon making IP and TCP mechanisms not needed in an intra-LAN operation as options, but that all options had to be implemented. This meant that all LAN implementations need contain all the capabilities of the standard long haul versions, but the full range of capabilities was not necessarily invoked for every intra-LAN packet of data.

It is believed [45, 48] that significant savings in protocol header overhead and software processing can be obtained with the use of the "discretionary TCP/IP" and that this protocol interworks with standard TCP/IP networks with a simple translation at a gateway between the local and other network, thereby fulfilling all the requirements of a local network protocol.

The MITRE work [45, 48] also considered the use of error checksum in the TCP and IP packet header for intra-LAN traffic. It was concluded that the software-based TCP/IP checksum processing delays and error protection provided were not justified when employed with more powerful, better performing and hardware implemented ones in the basic underlying LAN transmission protocols. Where TCP/IP packets would traverse into the long haul networks, the standard TCP/IP checksums should then be invoked to provide end-to-end protections.

Overall, MITRE concluded that the "discretionary" version of TCP and IP did not offer commanding advantages but did recommend to include them in any forthcoming reviews of the TCP and IP standards. Further, that while the discretionary versions do reduce header overhead and cause no perturbation to the basic TCP and IP operations, this was considered as being short-sighted: a school of thought which considers making protocol selections based solely on header overhead and communication line utilization.

5.17.2 Simulation/Modeling of TCP/IP/Ethernet LAN

This subsection discusses effort on the LAN Study which was devoted towards development of a simulation/modeling capability for use in evaluating LAN-based networking protocols and the results obtained. The motivation for this was the recognition of the complexity involved in performance assessing multiple layers of protocols when configured in both the intra and inter-LAN topologies. An introduction is given, discussing the modeling objectives, followed by a description of the design approach. Results obtained to date are given. A detailed description of the simulation/modeling design is contained in [95], titled "Progressive Project Document - Local Area Network Inter Operability Study Simulation."

5.17.2.1 Objectives

A paper published by Didic and Wolfinger [58] formed the basis of the architectural approach to representing a LAN-based suite of protocols for networking. Of particular interest was the methodology employed by Didic and Wolfinger in representing the OSI layered protocol model and accounting for use of internal computer resources to implement the protocol-functions. Their paper was

the first identified to attempt to follow the OSI model so closely. In their paper [58], the following points were made, which the LAN Study found particularly relevant to the investigation being conducted:

"One of the areas requiring the use of simulation models is the integrated analysis of a hierarchy of communication protocols."

"In our opinion, designers of simulation systems should pursue as a final goal a modeling tool which comprises components comparable to those defined by the ISO Reference Model. It should enable its user to tailor his simulation system by configuring network nodes, layers, protocols and their attributes."

Our overall objective in the LAN Study has been to develop an architectural model for a LAN-based computer network initially and later to expand this into an internet of interconnected LAN's, using representations of wide area net effects. It was felt the model should allow the following:

- Description of communication protocols
- Allocation of resources within the communicating computers
- Generation of requests (workload) created in a real system by the users of the network

Simulation modeling is a recognized technique for the analysis of complex systems. By constructing a model incorporating system characteristics and reproducing its behavior over time, system performance under various conditions can be assessed. Since a simulation model captures the time-dependent aspects of both system functions and loading, complex total system behavior can be studied.

A discrete event simulation model was developed to assess the performance of integrated layered protocols in a local area network environment. The performance was quantitatively evaluated during experimentation by the collection of these statistics for each protocol and for the entire network:

- Throughput
- Delay time
- Queue lengths
- Retransmissions
- Resource utilization

Another objective was to provide maximum flexibility in the configuration of experiments. The user was permitted to select which protocols were to be modeled as well as to specify the following parameters:

- Protocol parameters, i.e., processing times and header sizes
- System configuration, i.e., number of nodes and rate of channel
- Traffic characteristics, i.e., message lengths and arrival rates
- Node processor characteristics, i.e., speed and other loading

5.17.2.2 Simulation Approach

A simulation model is an abstraction of the actual system it represents. It does not propose to be an exact replica where every minute detail is emulated. However, the time-critical aspects of the system can be modeled in great detail. The task of modeling involves the identification of system features relevant to the study goals and characterizing them in sufficient detail to meet these goals. A model, therefore, represents the relevant aspects of the system and studies the behavior of the system over time with respect to these factors.

5.17.2.3 Simulation Methodology

The first step in using simulation techniques to accomplish the goals of a modeling effort is to specify the requirements. During frequent user/analyst meetings it was determined that the requirements were:

- To allow the variation of protocol parameters and configuration.
- To simulate loading on the channel during a simulation experiment in order to provide realistic contention and utilization of the channel.
- To model in detail the processing required to handle the selected protocols as well as representing the other processing required to handle user applications. The simulation of processing required for protocols had to include task generators and protocol submodels for a variety of media-access and host-to-host protocols.
- To accumulate throughput, delay and queue statistics and report the results.

The second step involves designing a model which represents the system under study and provides the detail necessary to satisfy the objectives. The design was documented in a simulation-oriented pseudo code and was reviewed to assure its accuracy. The areas where it was decided that great detail would be used were the contention for the channel in the Ethernet protocol and the TCP protocol handshakes and state changes.

The model structure is described in the overview provided in the next paragraph section.

The third step is to convert the design to computer executable code. The model was implemented in SLAM II.5 (Simulation Language for Alternative Modeling) and the user interface was coded in FORTRAN 77. SLAM provided the high level simulation language constructs useful for the model while FORTRAN provided the capability to have an interactive user interface.

The model was implemented in a top-down fashion with great care taken that the framework be very modular to allow the later addition of other protocols and layers. Also, the model was very parameterized to provide the user great flexibility in configuring the system for an experiment.

The fourth step is to verify and validate the model. Verification is the process of determining that a model executes as intended. Techniques used for verification included traces, examination of the summary report, and structured walkthroughs.

Each protocol submodel was first tested with only one message entity, then with multiple entities. Each submodel was also tested separately before being combined and tested with other submodels.

Event traces were used to verify each submodel. They provided a step by step report of the location of an entity and the simulated time at every point in the program where simulated time passes. The various conditions to be tested were forced in order to trace that entities followed the correct routes through the network, that the correct amount of simulated time passed, and that entities had the correct attributes.

The summary report also provided other data used in debugging the program. Data was provided on the minimum, maximum, and current status of resources, files, and activities. Additional debugging information was available from the statistics collected on collisions, retries, interarrival times, message lengths, time in layers and contents of the TCP Transmission Control Block.

A structured walkthrough at which several programmers stepped through the implementation and evaluated the flow of logic was used to verify all submodels.

Validation is the process of determining that the model is an accurate representation of the actual system. The validation techniques used were hand calculations, comparison with published results and face validity checks with experts.

For example, the expected performance of the model could be hand calculated for the case where only one node is transmitting. The performance of the model was also compared to published results from other studies. The Ethernet submodel, for example, was compared with results released by the IEEE Project 802 Traffic Handling Characteristics Committee Report [52] and the TCP submodel was tuned to results reported by MITRE [45-48].

The fifth step was experimentation with the model. This involved the collection of data from many simulation model executions. This process is documented in Paragraph 5.17.2.5.

5.17.2.4 Simulation Model Structure

The LAN Interoperability Study simulation provides LAN model configuration, simulation, and reporting of statistics.

The subdivisions of the software are:

- Setup

The Setup submodel interfaces with the user to allow selection of network configuration, loading and protocols configuration.

- Model

The Generators submodels generate requests for the highest layer protocol being simulated. This models the requests from the client above the highest layer protocol in the nodes modeled in detail.

The Ethernet Generator submodel also may be used to generate loading on the channel to represent contention caused by nodes not modeled in detail.

The Protocols submodels simulate the passage of messages down through the protocol layers, across the channel, and back up the layers to the client.

- Reports

The Reports submodel produces a summary of the statistics collected in the other submodels.

Figure 5.17.2.4-1 illustrates the division of the software into submodels. A Control Flow diagram is shown in Figure 5.17.2.4-2.

- a. Setup

The Setup submodel interfaces with the user to allow selection of network configuration parameters and node configuration parameters. The Setup submodel is divided into several modules:

- Exec

The Exec submodel interfaces with the user via menu to allow selection of the secondary menus listed below and also forces the user to specify which protocols are to be modeled.

- General Setup

The General submodel interfaces with the user to allow selection of the simulation run time and node configuration parameters for six types of nodes. These parameters include the arrival rate and length of messages as well as processor and TCP connection characteristics.

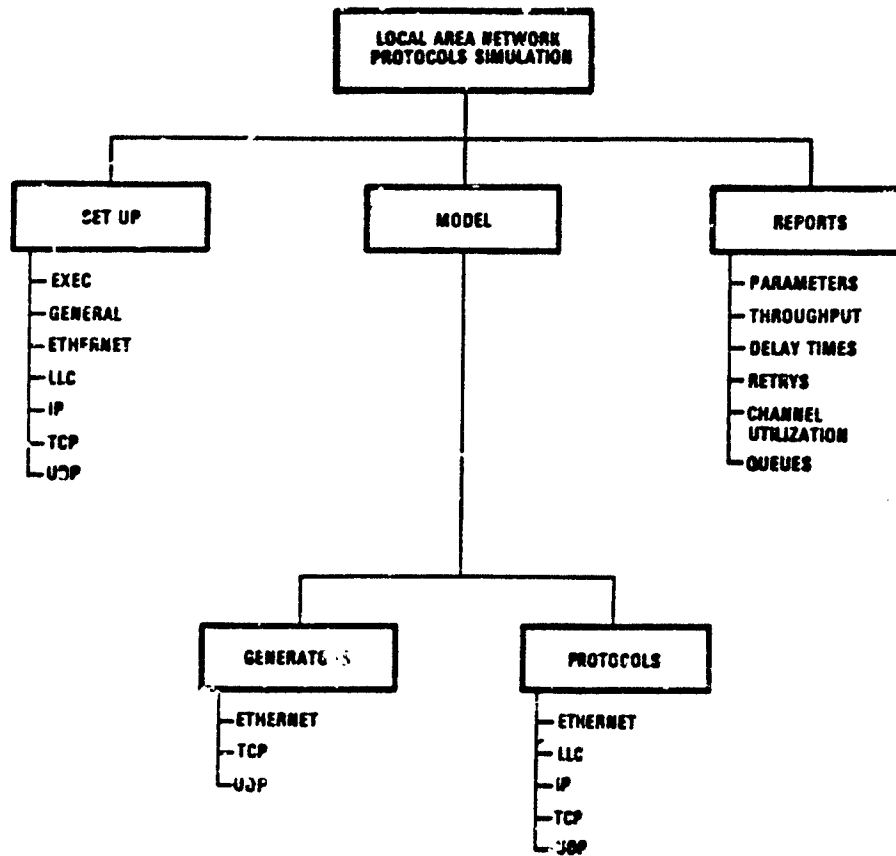


Figure 5.17.2.4-1. Model Partitioning Diagram

13457-31

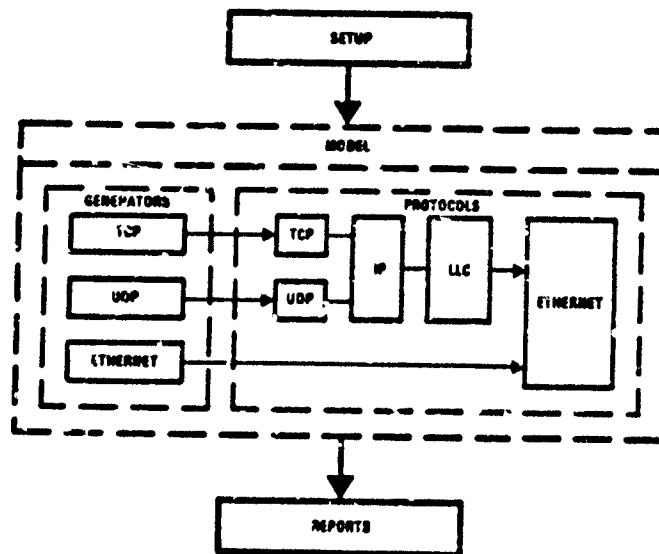


Figure 5.17.2.4-2. Control Flow Diagram

13457-32

- Ethernet Setup

The Ethernet Setup submodel interfaces with the user to allow selection of Ethernet parameters such as channel rate, number of jam bits, one-way propagation time, encapsulation time, and hardware interface to the host processor.

- LLC Setup

The LLC Setup submodel interfaces with the user to allow the selection of the logical link control parameter for processing time.

- IP Setup

The IP Setup submodel interfaces with the user to allow the selection of the internet protocol parameter for processing time.

- TCP Setup

The TCP Setup submodel interfaces with the user to allow the selection of the transmission control protocol parameters such as timeout period, time to process each state change, and transmission failure loss probability.

- UDP Setup

The UDP Setup submodel interfaces with the user to allow the selection of the user datagram protocol parameters for processing time and quantity of each type of node.

b. Model

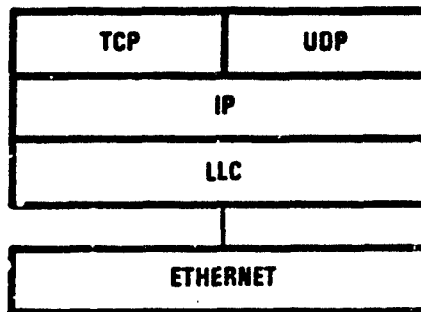
The Model submodels represent the arrival of requests for message transmission from the client layer and represent the processing required by each protocol selected. The model is subdivided into these submodels:

- Generators

The Generator submodels create entities to be passed to the protocol submodels. The generators represent the client layer above the highest protocol layer being modeled and loading at the media-access level.

- Protocols

The Protocol submodels simulate the operations of the selected protocols. Message entities are passed from one protocol submodel to another as specified in Setup. Figure 5.17.2.4-3 shows the protocol submodels currently implemented.



13457-33

Figure 5.17.2.4-3. Implemented Protocols Diagram

5.17.2.5 Results - Simulation Experiments

The model was designed and implemented to be very parameterized so that many experiments could be easily configured to study the effects of varying specific parameters. The results of these experiments are summarized in the following pages.

The protocols of greatest interest during this study were Ethernet and TCP. They were modeled in more detail and more experimentation was done to determine their performance than with the other protocols as resources to be devoted to this were constrained to a subset of the overall objectives set out at the start of the study.

5.17.2.5.1 Ethernet Experimentation Results

The performance characteristics of interest for the Ethernet protocol were throughput, one-way delay, and number of collisions as a function of channel (medium) loading utilization, number of nodes transmitting and message length. The results obtained are given in Figures 5.17.2.5.1-1 through 5.17.2.5.1-4.

In all the Ethernet experimentation, the model was configured per the following Ethernet blue book parameters [96]:

Number of jam bits	32
Number of preamble bits	64
Number of framing bits	144
Interframe delay time	9.6 microseconds
Backoff slot time	51.2 microseconds
Channel rate per second	10 megabits

The following estimated processing times were used:

Transmit encapsulation time	500 microseconds
Receive decapsulation time	1000 microseconds

Figure 5.17.2.4-3 shows the Ethernet portion in relation to the higher layer protocols (LLC, IP, TCP, UDP). The results obtained for the Ethernet were done alone, without any higher layer protocols implemented. Instead, they were represented as "client" loadings and presented traffic to Ethernet to process. In the figures which give the results (Figures 5.17.2.5.1-1 through 5.17.2.5.1-4), Ethernet All Data refers to total traffic on the cable medium, Ethernet Client Data refers to external representation of the higher layer protocols, message lengths are in bits and the nodes which loaded the cable medium were simulated by traffic generators.

Ethernet All Data Throughput (Figure 5.17.2.5.1-1)

Several simulation runs were made where four different sized client data packets were utilized (500, 1000, 2000 and 4000 bits per packet) while loading on the cable (medium) was increased at discrete amounts (1, 10, 50 and 100 nodes). Each node was caused to always have client data for transmission. As a result, this caused collisions on the cable to increase as more and more nodes were added. The results demonstrated two properties, as follows:

1. Total potential cable throughput decreased as a function of increasing nodes (collisions) on the cable
2. Increased sized packets yielded better throughput performance (less cable accessing occurs with larger sized packets)

As the number of nodes was varied from 1 to 100, the cable throughput potential was reduced by more than 50 percent (for 1000 bit packet size this changed from approximately 9 down to 3 Mb/s). Therefore, while a 10-Mb/s Ethernet cable system might be being employed, the actual potential system throughput at the cable is very sensitive and is determined by the total demand for bandwidth from those nodes attached.

Ethernet Client Data Throughput (Figure 5.17.2.5.1-2)

This shows the corresponding throughput values that were achieved at the client interface to Ethernet (client is the representation of Ethernet's higher layers). The results were the same as that which occurred at the cable (medium) except for a decrease in throughput overall caused by the Ethernet overhead. Therefore, the client throughput values are less than the All Data values.

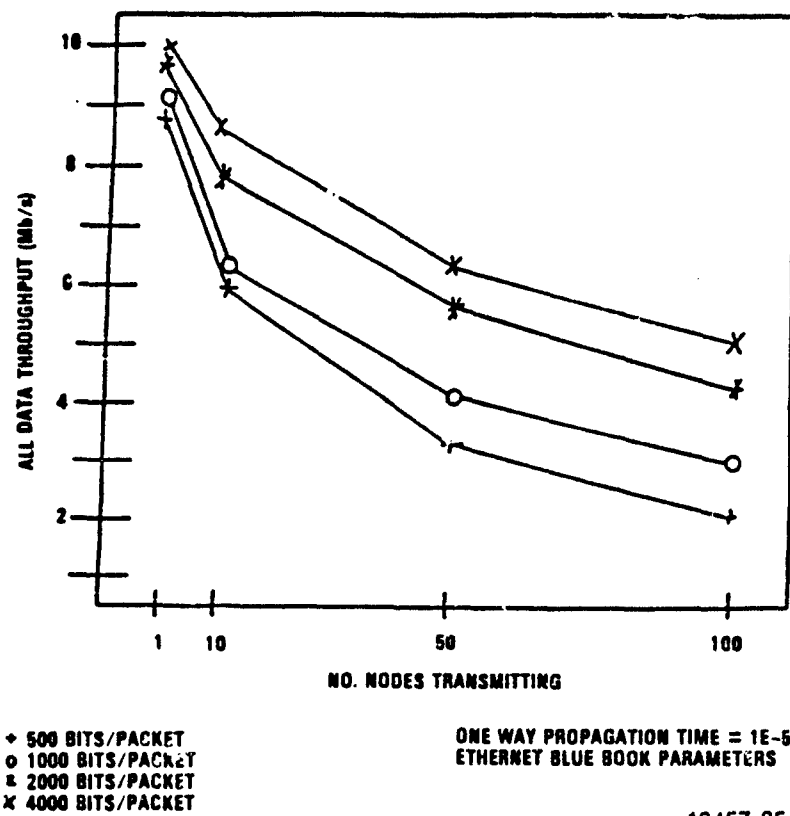
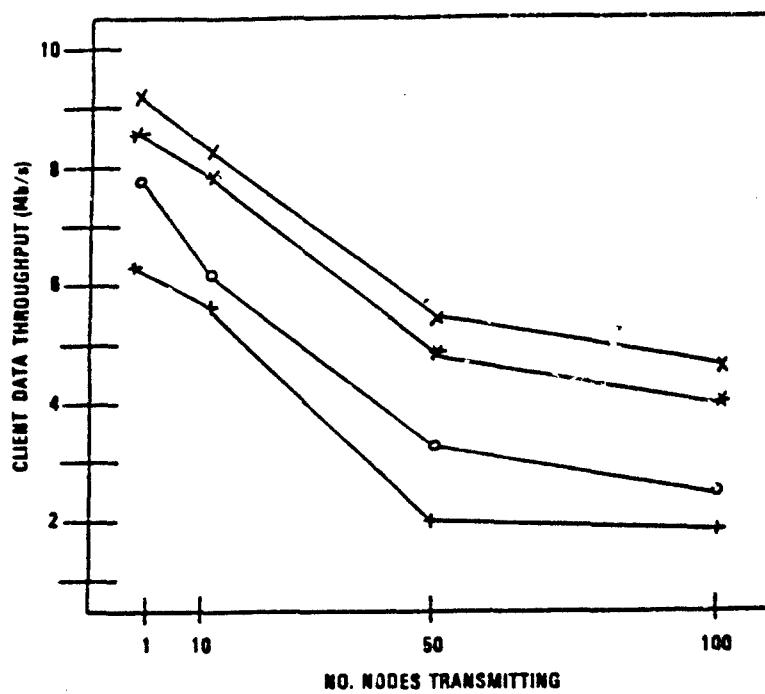


Figure 5.17.2.5.1-1. Ethernet All Data Throughput

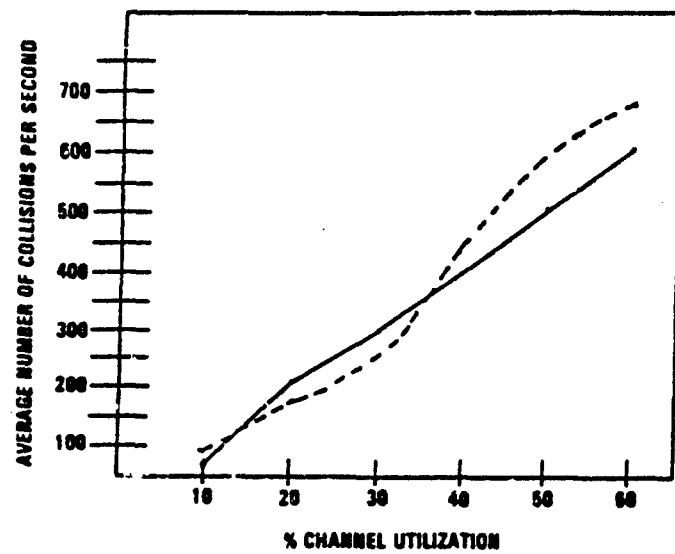


HTS/PACKET
 BITS/PACKET
 BITS/PACKET
 BITS/PACKET

ONE WAY PROPAGATION TIME = $1E-6$
 ETHERNET BLUE BOOK PARAMETERS

13457-36

Figure 5.17.2.5.1-2. Ethernet Client Data Throughput

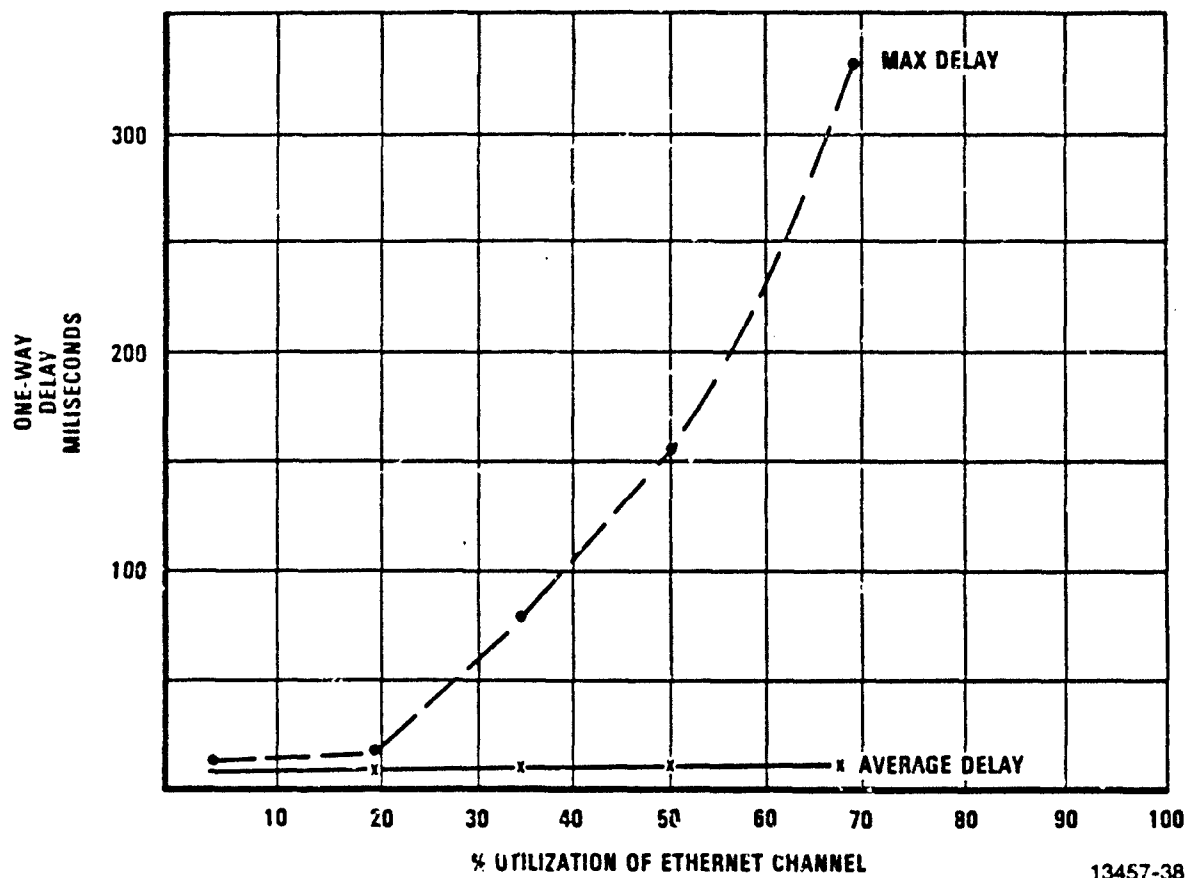


MESSAGE LENGTHS: 8592
 ONE-WAY PROPAGATION TIME: $1E-6$
 ETHERNET BLUE BOOK PARAMETERS

10 NODES: ———
 100 NODES: - - - -

13457-37

Figure 5.17.2.5.1-3. Ethernet Collisions



13457-38

Figure 5.17.2.5.1-4. Ethernet One-Way Delay

Ethernet Collisions (Figure 5.17.2.5.1-3)

While simulation runs were being made data was collected on the number of total Ethernet collisions which occurred at the cable level. The figure shows this for the case of 10 and 100 nodes as a function of % channel utilization of the potential 10 Mb/s capacity. The results indicated that collisions increased almost linearly as the loading on the cable increased for the 10-node case, but not as linear for the 100-node case. This difference was believed caused by the different run lengths in time that were employed. Further, the Ethernet collisions seemed to be independent of the number of nodes actively impressing the loading on the channel. The increase in channel utilization was seen to be the major contribution to throughput reduction seen in the previous figures and produced the decrease as the number of collisions which occurred increased.

Ethernet One-Way Delay (Figure 5.17.2.5.1-4)

For some applications, one-way delay is a more critical performance criteria than is throughput. Measurements were made of Ethernet. Results were obtained for the average and maximum one-way delays from source Ethernet Client to destination Ethernet Client.

The results show that average delay remained nearly constant across the channel loadings employed but that maximum delay exhibited a much more severe degradation as 60-65 percent was incurred in loading. This follows the theoretical prediction of delay-throughput for CSMA/CD systems.

The throughput and one-way delay values shown on these graphs are the maximum values achieved in a finite number of runs in which the interarrival rates of messages were varied.

5.17.2.5.2 TCP and IP Experimentation Results

The performance characteristics of interest for the TCP protocol were throughput, one-way delay, and maximum queue length as a function of input rate, CPU speed factor, and channel utilization.

For the TCP experimentation the processing times were configured as follows. These numbers were derived from statistics published by MITRE [45-48] on the time for TCP to process data. Other processing times were estimated from MITRE furnished design data for TCP by comparison of number of lines and complexity of code with the process data time. These comparisons were made using a C programming language implementation of TCP.

Attach header processing time	0.015 ms	Deliver data processing time	9.0 ms
Active open processing time	5.4 ms	CRC compute time per bit	0.23 ms
Send data processing time	5.9 ms	Close response processing time	3.8 ms
Close processing time	1.1 ms	Close success processing time	5.6 ms
Open response processing time	5.6 ms	LLI receive processing time	0.01 ms
Open success processing time	4.7 ms	ULI receive processing time	9.0 ms
Establish processing time	3.2 ms	TCP receive processing time	0.01 ms

For IP the following values were used:

IP Send	2.5 ms
IP Receive	2.5 ms

A value of 2 ms was used for the IP to Ethernet Interface Time.

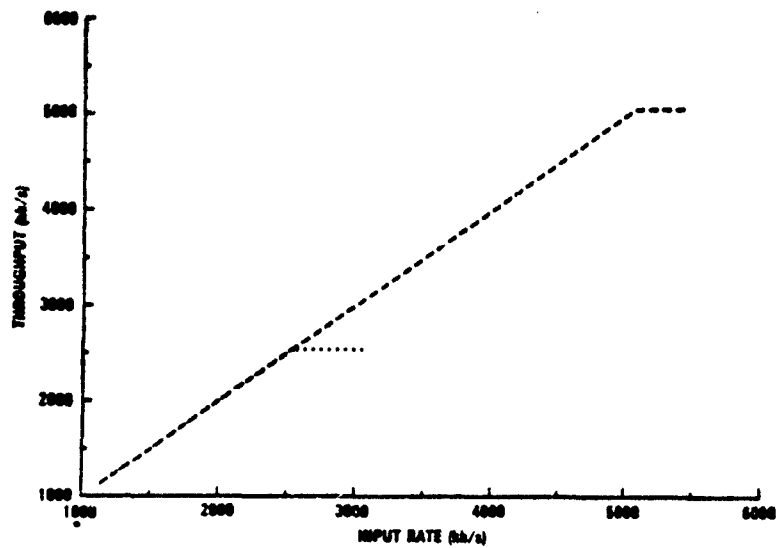
Input Versus Throughput Rate for Single TCP Connection-Single Node
(Figures 5.17.2.5.2-1 and 5.17.2.5.2-2)

TCP and IP protocol processing times were added onto the previously discussed Ethernet models. Most of the TCP protocol functions were represented, except for the flow control and window management operations (there was not enough time to investigate their effects). The objective was to determine the potential best performance obtainable in a LAN environment and sensitivities to different representations of internal machine processing speeds. At the CPU speed factor given, protocol processing occurred for the most part as if the CPU was always available. The model was designed to enable representing having to wait to be served by the CPU but all the results obtained did not employ that characteristic. For the cases where the CPU Speed Factor was changed, the processing times per function were reduced by a factor of 2 for each change in CPU Speed Factor (i.e., 1.0 was the slowest CPU, 0.5 was twice as fast, etc).

The results in the two figures indicated a potential capability of achieving throughputs for long sized TCP messages (11,440 bits) ranging from approximately 650 kb/s (Speed Factor of 1.0) up to 5.2 Mb/s (Speed Factor of 0.125). These results were obtained for the case of a single transmitting node with one open TCP connection and no other traffic loadings on the Ethernet cable (no collisions). This represented the theoretical best obtainable.

In each condition measured, the input-output results remained linear up to a point and then saturation occurred where for more input the corresponding output would not continue to increase.

TCP AND IP RUNS
(USER MESSAGE SIZE = 11,460 BITS)

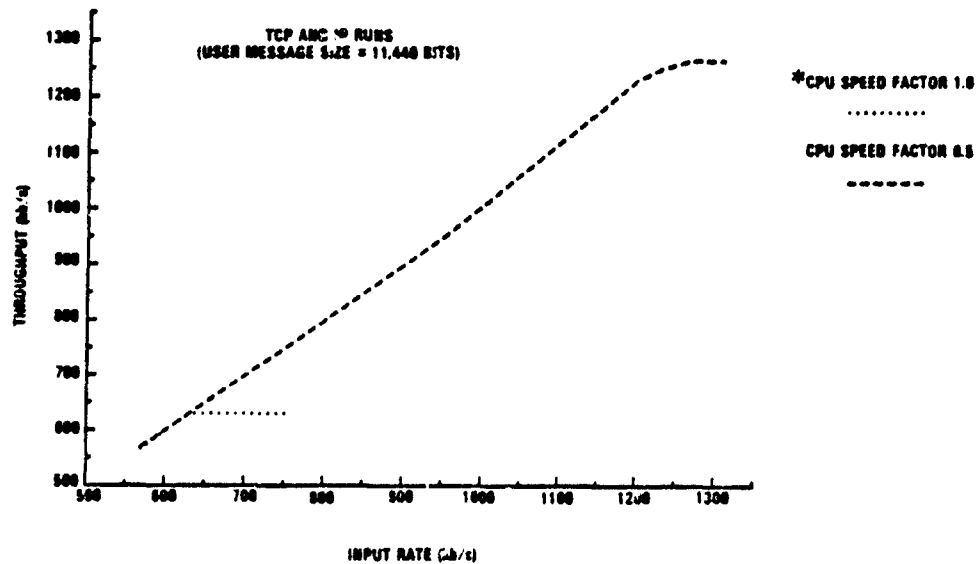


*CPU SPEED FACTOR IS A MULTIPLIER OF THE TCP TO
TCP 1-WAY PROCESSING TIME.
(INITIAL VALUE = 28 mSEC)

..... *CPU SPEED FACTOR = 0.25
----- CPU SPEED FACTOR = 0.125

13457-30

Figure 5.17.2.5.2-1. Input Versus Throughput Rate for CPU Factors 1, 0.5 (Single TCP Connection-Single Node)

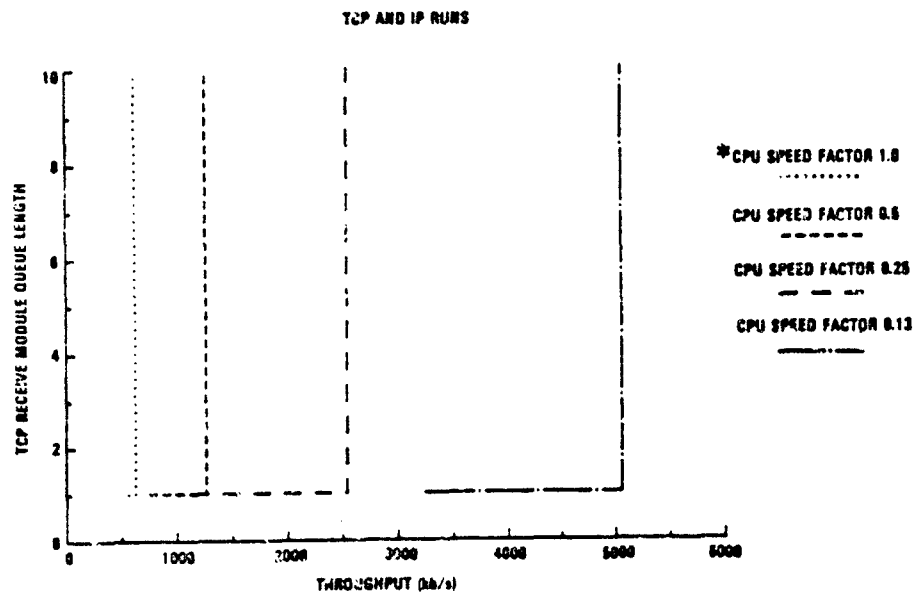


*CPU SPEED FACTOR IS A MULTIPLIER OF THE TCP TO
TCP 1-WAY PROCESSING TIME.
(INITIAL VALUE = 28 mSEC)

*CPU SPEED FACTOR 1.0
.....
CPU SPEED FACTOR 0.5

13457-40

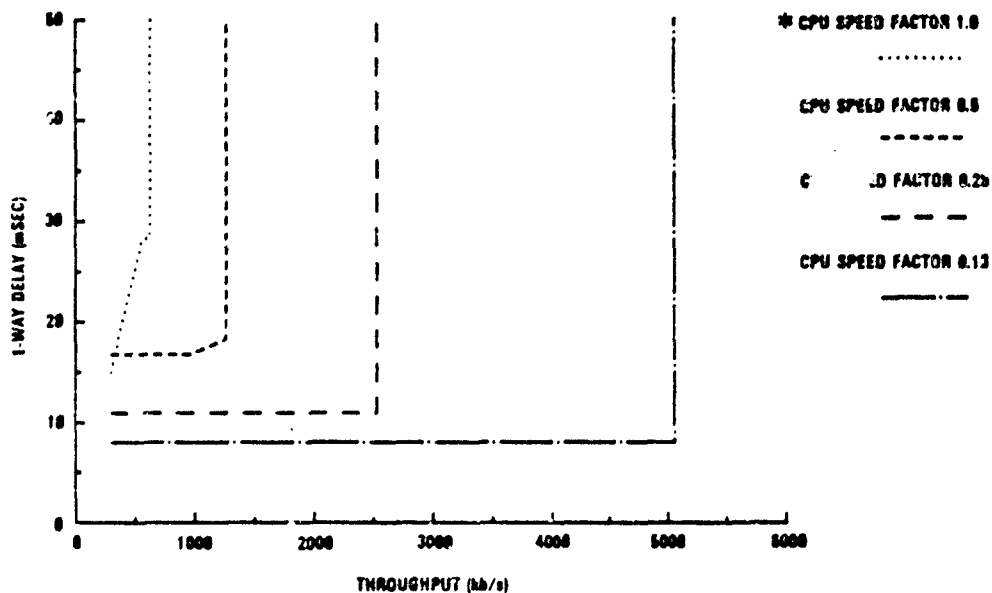
Figure 5.17.2.5.2-2. Input Versus Throughput Rate for CPU Factors 0.25, 0.125 (Single TCP Connection-Single Node)



* CPU SPEED FACTOR IS A MULTIPLIER OF THE TCP TO
TCP 1-WAY PROCESSING TIME.
(INITIAL VALUE = 25 mSEC)

13457-41

Figure 5.17.2.5.2-3. Throughput Rate Versus Max Queue Length
(Single TCP Connection - Single Mode)



* CPU SPEED FACTOR IS A MULTIPLIER OF THE TCP TO
TCP 1-WAY PROCESSING TIME.
(INITIAL VALUE = 25 mSEC)

13457-42

Figure 5.17.2.5.2-4. Throughput Rate Versus One-Way Delay
Time (Single TCP Connection - Single Mode)

Throughput Rate Versus Max Queue Length for Single TCP Connection -
Single Mode (Figure 5.17.2.5.2-3)

As the model was run in the previous examples for TCP, data was gathered on the TCP Receive Module Queue Length. The results showed a constant value of approximately 1 for a range on increasing throughput up to a point, where thereafter the size of the queue increased an order of magnitude. It was concluded this condition is what produced the throughput saturation effects for the four cases previously shown.

Throughput Rate Versus One-Way Delay Time for Single TCP Connection -
Single Node (Figure 5.17.2.5.2-4)

One-way delay time data was obtained while the CPU Speed Factor was changed from 1.0 through 0.125. The results showed that faster CPU speed (corresponding reduced protocol processing time) improved performance by reducing the one-way delay. For each condition, there was a range of operation where delay was minimal and constant up to a point in throughput where, upon reaching saturation, the delay increased to very large values.

CPU Speed Factor Versus Throughput for Single TCP Connection - Single
Node (Figure 5.17.2.5.2-5)

This figures show the maximum values of TCP throughput that were obtained at the saturation points for the four values of CPU Speed Factors.

Channel Utilization Versus Throughput for Multiple TCP Connections -
Multiple Nodes (Figure 5.17.2.5.2-6)

The model was reconfigured to create the case where three nodes were simulated, with single and multiple TCP connections established among them. In addition, loading nodes were added to cause the total Ethernet channel (medium) to be varied from approximately 5 to 83 percent loading (0.5 to 8.3 Mb/s).

Three different TCP message sizes (11,440 bits, 512 bits and 8 bits) were employed on different connections and operated on the Ethernet channel concurrently. Measurement results were obtained which indicated that over a large range of loading, throughput was not affected. Around the 65 percent loading region total system throughput began to decrease. The results for throughput were not set to indicate maximum attainable but representative of typical system operation under the condition of loading from other nodes.

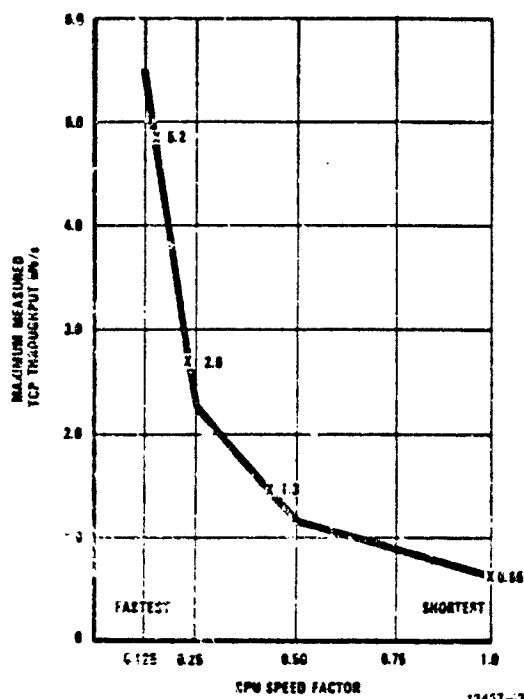


Figure 5.17.2.5.2-5. CPU Speed Factor Versus Throughput
(Single TCP Connection - Single Mode)

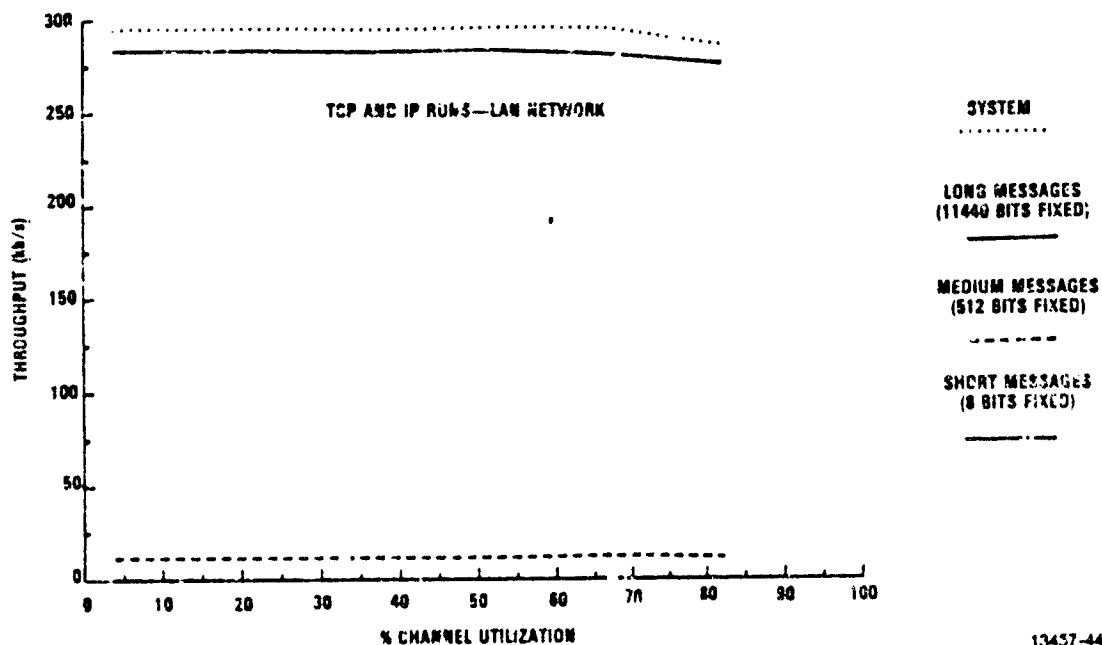


Figure 5.17.2.5.2-5. Channel Utilization Versus Throughput
(Multiple Connections - Multiple Access)

Channel Utilization Versus One-Way Delay Time for Multiple Connections -
Multiple Nodes for Fixed 11,440 bit, 512 bit and 8 bit Message Sizes
(Figures 5.17.2.5.2-7 through 5.17.2.5.2-9)

These figures show 1-way delay measurement results obtained for the previous three different message sizes for TCP throughput. The values for Ethernet and IP are the average delays whereas both the average and maximum delays are given for TCP. As was shown earlier, Ethernet average delay remained essentially flat whereas Ethernet maximum delay increased sharply as 60-65 percent and greater channel loadings were experienced.

The results indicate how this maximum Ethernet delay characteristic affects the protocol layers above it (TCP, IP). Up to around 50 percent channel utilization, delay was constant but as loading was increased, delays increased to large values, even though throughput remained constant up until about 65 percent.

5.17.2.5.3 Other Results

Network Minimum Header Overhead Pie Chart (Figure 5.17.2.5.3)

This figure shows the relationship of protocol headers (in percentage) to that of actual TCP Client or User data, for the message sizes of 8 bits, 512 bits and 10,240 bits. For the 8-bit case, data represents only 1 percent (overhead of 99 percent) while for 512 bits, data is 46 percent (overhead 54 percent) and for 10,240 bits, data is 94.5 percent (overhead of 5.5 percent).

Protocol Layered Data Rates (Table 5.17.2.5.3)

This table shows the throughput data rates as a function of protocol layer when TCP, IP, LLC and Ethernet were used together to produce a TCP Client rate of 272 kb/s for 11K bit message size.

5.18 TCP Alternatives for Intra-LAN Use

This reports on three references works [49, 50, 51] which examined TCP for use in LAN's where the predominate traffic would be intra-LAN oriented.

5.18.1 The Local Network Transmission Control Protocol (LNTCP) Alternative

Reference [49], by a member of the Naval Ocean Systems Center, discussed possible approaches to applying TCP within a LAN-based command center environment. It concluded that "because of their very different topology, transmission media, and other features, delay, throughput, and cost considerations indicate that long haul networks and local area networks should use very different host-to-host communication protocols. Long haul networks require complex protocols which efficiently utilize communication channel bandwidth at the expense of processing time. Local area networks require simple protocols which expend communication channel bandwidth in order to reduce processing time." The reference goes on further in considering its conclusions, as follows:

TCP AND IP RUNS—LAN NETWORK
LONG MESSAGES (11140 BITS FIXED)

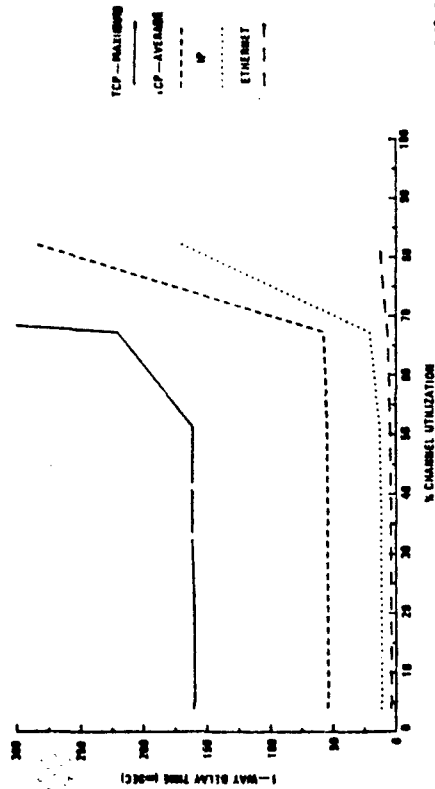
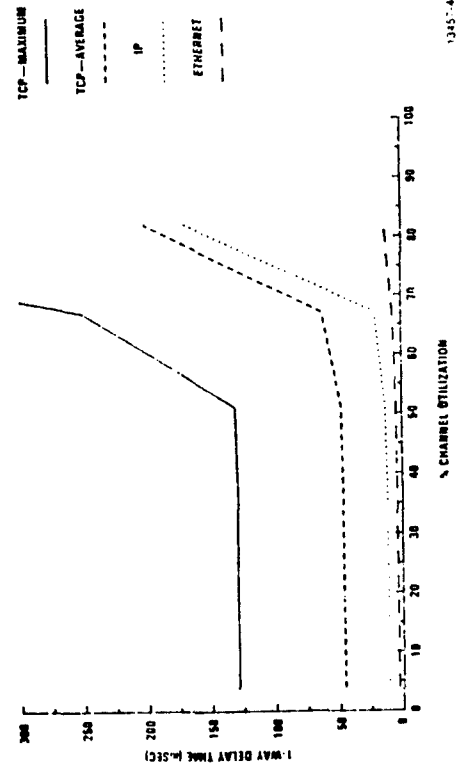


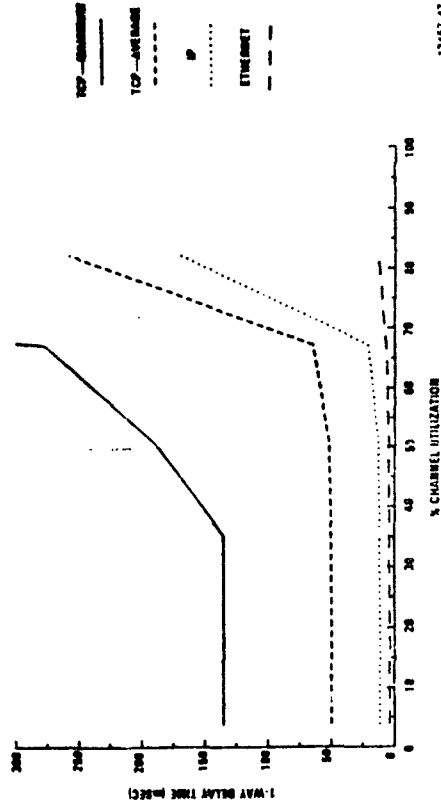
Figure 5.17.2.5.2-7. Channel Utilization Versus One-Way Delay Time (Multiple Connections - Multiple Modes, Fixed 11,140-Bit Message Size)

TCP AND IP RUNS—LAN NETWORK
MEDIUM MESSAGES (612 BITS FIXED)



5.17.2.5.2-8. Channel Utilization Versus One-Way Delay Time (Multiple Connections - Multiple Modes, Fixed 612-bit Message Size)

TCP AND IP RUNS—LAN NETWORK
SHORT MESSAGES (8 BITS FIXED)



5.17.2.5.2-9. Channel Utilization Versus One-Way Delay Time (Multiple Connections - Multiple Modes, Fixed 8-bit Message Size)

LOCAL AREA NETWORK
MINIMUM HEADER OVERHEAD

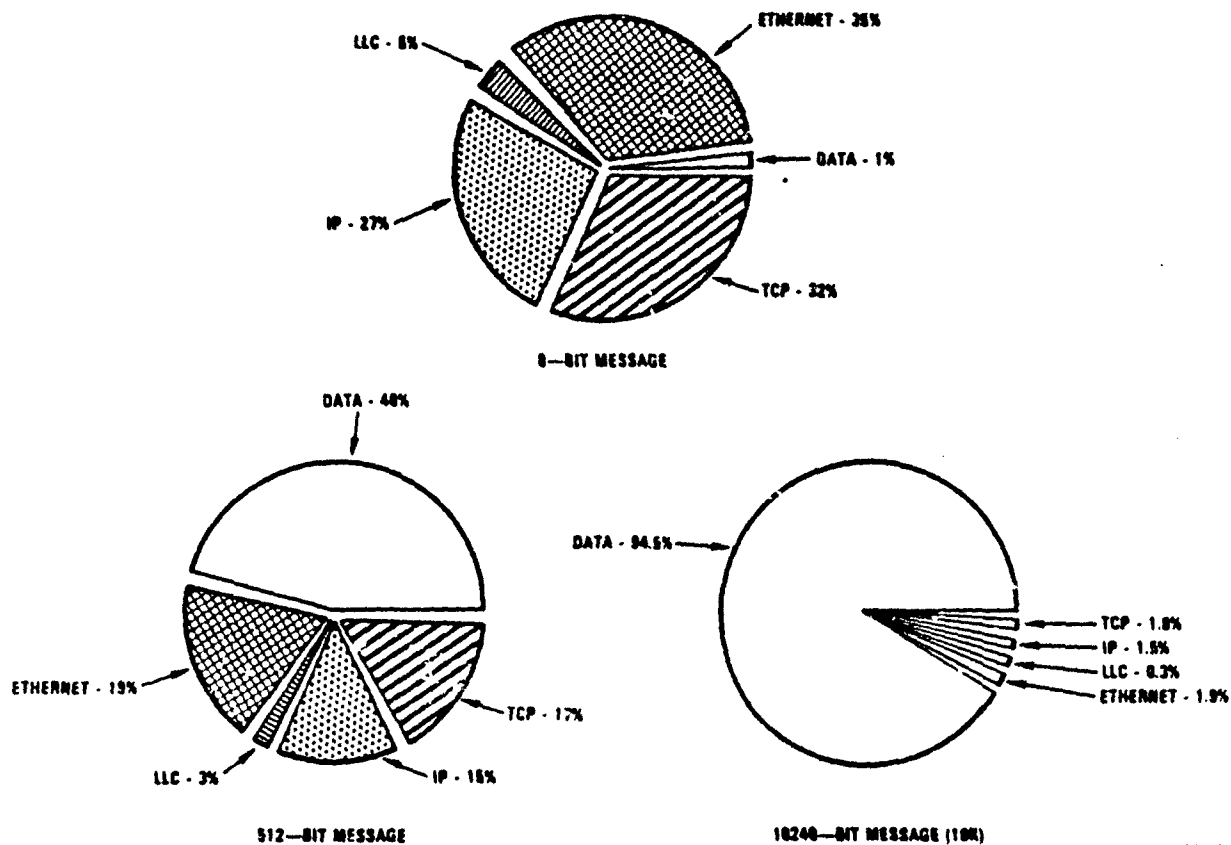


Figure 5.17.2.5.3. Network Minimum Header Overhead Pie Chart

13457-48

Table 5.17.2.5.3. Protocol Layered Data Rates

Channel Rate	10.0000000 Mb/s	10000.0000 kb/s
Ethernet All Data Rate	0.3000896 Mb/s	300.0896 kb/s
Ethernet Client Data Rate	0.2901819 Mb/s	290.1819 kb/s
LLC Client Data Rate	0.2886576 Mb/s	288.6576 kb/s
IP Client Data Rate	0.2810363 Mb/s	281.0363 kb/s
TCP Input Data Rate	0.2722720 Mb/s	272.2720 kb/s
TCP Client Data Rate	0.2718907 Mb/s	271.8907 kb/s
UDP Client Data Rate	0.0000000 Mb/s	0.0000 kb/s
Average Interxmit Time	0.0203467 sec	20.3467 msec

Conditions: 2 nodes and 1 TCP Connection

11k bits message size (fixed size)

Interarrival time of 41 ms (averaged)

"TCP is a complex protocol with features designed to efficiently utilize the communication facilities in a long haul network such as the ARPANET. Any local network interconnected with the ARPANET which uses TCP as its only host-to-host protocol will achieve relatively simple internetwork communications at the expense of definitely suboptimal intra-network communication performance."

It goes further:

"An alternate approach would be to implement a Local Network Transmission Control Protocol (LNTCP) which retains many of the features of TCP, but substitutes for the TCP flow control, resequencing, and duplicate detection features, and which uses a much larger packet size than the ARPANET. Such a protocol should provide for efficient local intra-network communication performance while allowing for good internetwork communication performance as well."

Lastly it states:

"Each Command Control Network node should be able to implement both the LNTCP and TCP with a limited amount of effort and cost. Being identical in many respects, the two protocols could share substantial portions of software. Those sections of software unique to either LNTCP or TCP could be either executed or bypassed, depending on the situation."

5.18.2 An Extended Backplane Approach to LAN Interprocess Communications

A paper by Cheriton [50] at Stanford University recently described a novel approach to LAN interprocess communications which departs from the use of standard TCP at almost the extreme. A distributed operating system, termed the V-System (pronounced "Vee"), is a system at Stanford which uses Ethernet. However, it uses the Ethernet LAN as an extended backplane to connect a collection of diskless SUN workstations and server machines. A distributed kernel provides uniform, transparent message-based interprocessor communication on a single workstation and between processes on different workstations.

The paper reports the following:

"We have a strong need and desire to support and use standard internetwork protocols, in particular, being part of the ARPA Internet community, the IP/TCP family of protocols. Normally, use of internetwork protocols is not a problem. However, our use of a local network as an extended backplane means that local network performance must be compared to intra-machine communication performance, not long-haul network performance. Using this comparison, we argue that using internetwork protocols on a local network for interprocessor communications leads to minimal benefits and significant costs."

The paper further discusses the characteristics of distributed processing and internetwork protocols as follows:

"Internetwork protocols do not provide a communication model well-suited to distributed systems. A distributed system requires a transport-level facility for interprogram and possibly intraprogram communication that works transparently within a single machine as well as across the network between machines. Because of the predominant use of procedure call-like communication within systems (even message-based systems) the appropriate model is some form of request-response transaction style communication, where a transaction typically consists of requesting a service and getting a response back."

In the concluding remarks, the following was given:

"The distributed V-System achieves high performance on the local network with a transparent network interprocessor communication facility based on a simple request-response interkernel protocol. Measurements indicate that this facility gives performance superior to internetwork protocols and the performance is sensitive to additional protocol overhead. In particular, processor time is the critical factor. The interprocessor communication facility suffices as the single transport-level communication mechanism on the local network as well as for all interprogram communication, whether local or distributed. It is as efficient to use for local network file access, file transfer and terminal access as more specialized protocols."

5.18.2 Protocol Functional Approach to LAN

A paper by Schneidewind [51], with the Naval Post Graduate School, discusses three approaches for solving the interconnection problem: network access, network service, and protocol functions. Of the three, the protocol method provides one set of protocol functions for the local network and another set for the long-distance network, where some functions, or corresponding network layers, are common to both networks.

When trying to optimize local network communications, as the primary user objective, the paper advocates "using only those layers and protocols necessary for local network operation, while also providing communication between local network via the long-distance network."

The paper states further:

"It is not necessary to implement all ISO layers in the local network to achieve effective intra-local network communication, nor is implementation necessary to connect to a long-distance network. However, the number, types, and characteristics of the layers utilized determine the efficiency of interlocal network communication (i.e., communication over the long-distance network)."

"The protocol functions approach solves the problem of local network communications efficiently, but at high hardware and software costs. Its use of the protocol translation in the gateway necessitates a complex network interface."

In an example approach cited, the following was given:

"The protocol functions approach provides only those protocols in the local area network that are needed to support this type of communications environment. The LAN has no need for the services provided by the transport and network layers, because routing, switching, and traditional flow control and congestion control services are unnecessary. The presentation layer, implemented in the terminal management module, will accept data from the application process and convert it to LAN format. Conversely, it will accept messages in the LAN format and convert them to the appropriate application process format." This is illustrated by Table 5.18.3 and Figure 5.18.3.

"To simplify the LAN design, the following message formats are used:

1. TCP format will be provided to the DDN by the NC module whenever communication on the DDN is necessary. A much simpler format will be used for the intra-LAN communication.
2. End-to-end virtual circuit connections and breaking of a complete message into fragments, services normally provided by the transport layer, will be implemented in each of the LAN modules. End-to-end in this context refers to the logical communication linkage between two modules separated by a relatively short distance; in some cases the two modules could be in the same hardware unit."

In closing, the paper stated: "Where optimization of intra-local network performance is desired, this is accomplished by using only those layers and protocols that are compatible with and can take advantage of the characteristics of local networks."

5.19 Generic Gateways for LAN Interoperability

5.19.1 Introduction

Local area networks (called LAN's) are the latest technology to offer a major breakthrough for building distributed information systems. LAN's provide high data rate (1-200 Mb/s) peer-to-peer digital communications, at low cost, for interconnecting system elements in a localized area (i.e., a room, building or a campus). Tradeoffs in throughput-delay performance, media and media access method, topology, connection-connectionless transfer service and methods for interconnecting LAN's, are possible when choosing a solution to a system's problem.

Table 5.18.3. Use of ISO Layers in LAN Design

<u>Layer</u>	<u>LAN Communication Protocol/Module</u>	<u>DON Communication</u>
Application	Application Process Modules	Same as for LAN
Presentation	Terminal Management	Terminal Management
Session	Session Services	Session Services
Transport	--	TCP
Network	--	IP
Data Link	Local Communications	Specified by the DON (Various Protocols)
Physical	Local Communications	

This discusses issues associated with interconnecting LAN's to other LAN's and to non-LAN's. This is discussed in the context of the emerging international consensus for how to interconnect open systems, the need to understand the IEEE 802 family of LAN protocols, the principles which interconnection is based upon, a generic family of gateway interconnection elements and example interconnection situations. The beneficial impact of understanding the issues presented herein can be realized in the economic savings and performance gain achieved through the application of these principles to real systems problems.

This material is based upon reference [60] which deals with generic gateways usable to interconnect LAN-based heterogeneous end system or system elements.

5.19.2 Objectives

This provides insight to the designers, developers and users of mixed media-based LAN's, into the bigger systems picture of which LAN's are a small but very important part. Sometimes in the quest to apply better performing LAN technology for achieving higher and higher data rates, it is easy to not see how this capability is to fit into the overall system. In the context of this discussion, the bigger picture is one where systems elements (computers, storage, processing, users, resources) will be physically/logically distributed, interconnected together by an interprocess communications subsystem employing LAN

P — PROCESS: PROGRAM IN EXECUTION
 FM — FUNCTIONAL MODULE (E.G. TERMINAL MANAGEMENT)
 NC — NATIONAL COMMUNICATIONS MODULE: PROTOCOL CONVERTER
 SS — SESSION SERVICES MODULE: SESSION LAYER PROTOCOL
 TCP — TRANSMISSION CONTROL PROTOCOL: TRANSPORT LAYER PROTOCOL

IP — INTERNET PROTOCOL: INTERNET LAYER PROTOCOL
 HIP — HOST-IMP PROTOCOL: NETWORK LAYER PROTOCOL
 IIP — IMP-IMP PROTOCOL: NETWORK LAYER PROTOCOL
 ☑ — PROTOCOL CONVERSION

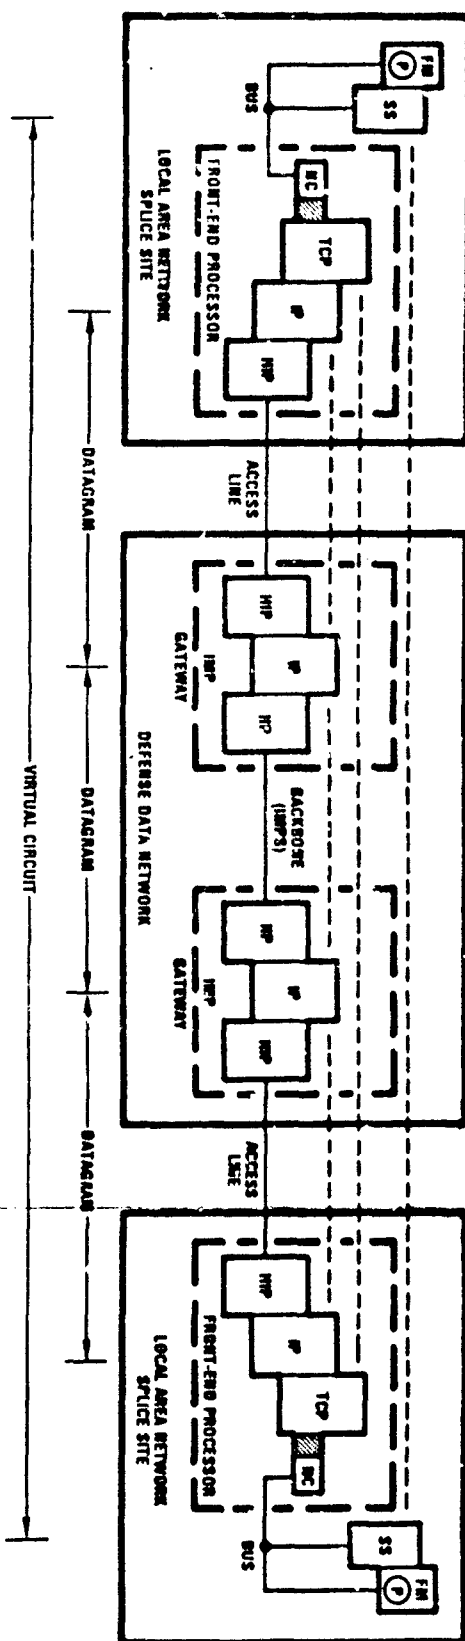


Figure 5.16.3. Relationships Between a Local Area Network and the Defense Data Network

technology, and then functionally integrated to interoperate as a virtual global system to the users. A global network operating system will manage the allocation and use of system resources.

5.19.3 Open Systems Interconnection

International consensus has now been achieved on the way to interconnect systems in an open (nonproprietary) way [5]. This agreement, reached in 1984 by both the ISO (International Standards Organization) and the CCITT (International Consultative Committee of Telegraphy and Telephony) are documented in ISO 7498 and CCITT X.200 documents for Open Systems Interconnection (OSI).

The services provided by the seven layers of the OSI/RM can be divided into three main regions, as follows:

Networking Utilities - Layers 7, 6, 5

Host-to-Host Transport and Internetworking - Layers 4, 3

Subnetwork Transmission (Local and Wide Area Networks) - Layers 3, 2, 1 and Media 0

The technologies associated with LAN's fall within the subnetwork transmission region to form one of the major emerging technology foundations for supporting the ability to build truly distributed processing systems. Public and private wide area networks (WAN's) provide a range of alternatives in selection of design solutions for layers 3-0 of the OSI/RM. Overall responsibility for end-to-end data transport reliability of packets fall to the protocols which implement layers 4 and 3. Finally, protocols of layers 7, 6, and 5 respectively deal with the semantics, syntax and organization of information exchanges. In the highest region of the OSI/RM, a set of generic (canonical) virtual protocols provide a set of commonly used networking utilities to the ultimate end system user. Examples of networking utilities are virtual file, terminal, job, message, document and resource management [17].

5.19.4 LAN Architecture and Protocols

To date, the most widely documented and supported consensus for the architecture and protocols for LAN's is represented by the family of standards developed by the IEEE Computer Society's Project 802 [41]. There, three methods for accessing media were developed based upon CSMA/CD, Token Bus and Token Ring operations. A common agreement was reached for the next level of protocol called Logical Link Control (LLC). Together, these span the lowest region of the OSI/RM.

The CSMA/CD and Token Bus protocols were developed to operate with a broadcast bus topology while the Token Ring was for point-to-point topology. Coax cable was the medium specified for the bus while twisted pair and fiber-optic cable were intended for the point-to-point cases. Nevertheless, fiber-optic cable technology employing a star arranged topology is capable of supporting the Token Bus medium access control protocol.

5.19.5 Connectionless and Connection-Oriented Services [61]

In order to interconnect elements of the 802 architecture, a closer understanding of the services provided by the LLC and MAC entities is necessary. First, both the MAC and LLC provide a connectionless service, sometimes called datagram. However, an optional connection-oriented service may be provided by LLC in addition to connectionless. With the connectionless service, a single data frame at a time comprises the span of communications control. That is, one MAC entity in a station transmits a single MAC frame to one or more other stations by way of the interconnected media. The receiving station(s) process that entity as a stand-alone message. If successful (having no detected error and addressed to that station), then the single message is accepted and its control field(s) contents acted upon, else the frame is discarded.

Recovery from discarded frames may be performed by the MAC sending station, in the case of Token Bus and Token-Ring and/or a higher protocol. In the CSMA/CD case, lost frames are not recovered by the sending station's MAC entity. A higher layer protocol entity would, upon timing out, attempt a retransmission. When LLC employs the connection-oriented service, the sending LLC entity initiates recovery for CSMA/CD, Token Bus and Token Ring for non-MAC frames either triggered by a time-out event or based upon information sent to it from a destination station indicating the loss of a sequenced identified LLC message frame (MAC information field). When LLC employs the connectionless service, it does not perform any recovery operation in the event of a lost frame. A higher layer (transport) may perform this recovery.

A significant distinction can be made between these two types of service. When connectionless service is employed, each message frame is the entire entity in space over which control spans, whereas with connection-oriented service there is a time orientation during which control is spanned. With connectionless service, sending and receiving stations maintain minimum (or no) state information, but with connection-oriented there is both state and control block descriptive data managed at both ends of the established connection.

Connection-oriented service employs connection open, data transfer and connection close phases, while connectionless is always in the data transfer phase. These differences in operational characteristics hold at any layer in the OSI/RM where either or both forms of service are performed. Connection-oriented service in addition performs sequencing, flow control and error recovery operations which connectionless does not perform.

5.19.6 Principles of Interconnection

LAN and WAN networks, being offered by vendors and public carriers, in spite of the evolving standards work, are heterogeneous in many instances. Systems designers will be faced with having to interconnect them together in order to build an integrated global system. The main objective of interconnecting networks is the extension of communications facilities to a larger population of users.

This problem of network interconnecting was examined and reported on by Gien and Zimmermann [62]. In that paper, a set of principles is laid out for network interconnections based upon several simple but powerful structuring techniques. These techniques stress multiplexing, switching, cascading, wrapping, layering and the equivalence of services.

In order to interconnect networks, they must exhibit the following:

1. Levels of equivalent services, to be possibly merged into a global network offering these services
2. A set of properties which make interconnection viable, such as cascadability of services, multiplexed interfaces and interpretation of a global address space

If these constraints are not satisfied, the network(s) must be modified, usually by wrapping it (end-to-end) in an additional layer which externally exhibits the required interconnection capability.

Multiplexing

A mechanism is employed by multiplexing in either space or time to share a resource among several using entities.

Switching

Switching involves the interpretation of addresses and routing of requests when resources are shared.

Cascading

This consists of joining a series of entities in a linear string which forwards messages or propagates activities along the cascade. Cascading is the only way for communication between entities which are not directly connected to each other. This is like joining two similar service networks by plugging them together.

Equivalence of Services

Interconnection of networks can only be performed at a level providing equivalent services on either side [62, 63, 64]. In addition, these services must be cascable. Networks offering identical services, such as connectionless or connection-oriented, can be interconnected without great difficulty.

If this is not the case (equivalence of service), either a common subset of the service has to be chosen (subnetwork de-enhancement) or a new (subnet enhancement) sublayer has to be added on top of a "poor" subnetwork to achieve equivalent services [62].

Protocol Conversion

When two subnetworks are being interconnected which possess equivalent services (semantics) but have different formats, then protocol conversion must be performed [62]. Protocol conversion consists of expressing the semantics carried by one protocol in the message formats and protocol syntax of the other. In other words, it consists of converting the necessary message formats and protocol syntax to preserve the protocol semantics.

When the two architectures being interconnected carry protocol functions which differ in their distribution, then composite functions in two or more layers must be considered in order to find a point where the two services are equivalent. In this approach, protocol conversion can be extended to protocols pertaining to two (or more) layers, with the possibility of converting the protocol of level N of one architecture into the protocol of level M+1 of the other architecture. On the other hand, when two protocols belonging to different architectures support different sets of functions, protocol conversion can be applied effectively only to the functional subset common to both protocols. Lastly, when two protocols are not "convertible," the only possible solution for communicating with the other entity is actually to implement the other entity's protocol [67].

5.19.7 Gateways for Interoperability

A set of gateway elements, or building blocks, are needed in order to support full LAN interoperability in a heterogeneous world. This paragraph provides a summary of four generic types of gateways, spanning the full range of the OSI/RM. Appendix G presents a design approach discussion where each of these four is discussed in greater depth and an application approach presented.

Gateways are communications processing nodes used to interconnect networks employing heterogeneous protocols. Four types span the seven protocol layers:

- Bridge Gateway (through layer 2)
- Packet Switching Gateway (through layer 3)
- Protocol Conversion Gateway (layers 4-7)
- Open End Systems Interconnection Gateway (layers 1-7 between heterogeneous end systems)

Bridge Gateway (through layer 2)

This is employed to join two media access control LAN segments together; these may be homogeneous or heterogeneous. Logical Link Control using protocols is extended across joined physical LAN segments. When just two like media are joined this is called a repeater (baseband) or frequency translator (in the case of broadband). These can perform some of the address filtering/blocking in support of multilevel security for joined LAN's.

Packet Switching Gateway (through layer 3)

This is employed to interconnect LAN to LAN or LAN to WAN to LAN subnets to form an internet. Underlying layers of protocols may be homogeneous or heterogeneous. However, LAN logical links or WAN network virtual circuits are terminated at gateway. An end-to-end internet protocol, employed on top of joined LAN-WAN protocols, performs host-to-host packet switching, routing, fragmentation, reassembly.

Protocol Conversion Gateway (layers 4-7)

This is employed to interface two heterogeneous network system's upper layer protocols. This type converts from vendor A to vendor B specific protocol. Gateway functions are very vendor protocol and application specific and difficult to standardize.

Open End Systems Interconnection Gateway (layers 1-7 join heterogeneous end systems)

This type is employed to interconnect in an open manner (nonproprietary way) multiple vendor end systems, each employing their own internal heterogeneous protocol suites, to create a cooperative Global Network Operating System set of networking utility services.

This type spans layers 1-7 and provides a set of generic networking utility services, host-to-host transport, internetworking and LAN/WAN subnet services. This type provides services, functions and protocols necessary to support fully distributed data processing. It performs networking utility functions employed in building a distributed network operating system. The internal architectures of these end systems are not addressed by the Open System Interconnection Network Utility protocols. Two major architectures/protocol suites exist: DOD's and ISO's. The DOD/ISO proposal suites provide similar functionality for interprocess communication between application processes residing in heterogeneous or homogeneous end systems.

SECTION 6.0
CONCLUSIONS

6.0 CONCLUSIONS

6.1 General

This section presents a set of concise statements of the conclusions reached from the studies' and investigations' results given in Section 5.0. This is arranged in the same order of presentation as that given in Section 5.0.

6.2 Strategic Command, Control and Communications

1. C³ must survive and offer sufficient flexibility to identify, reconstitute, and employ surviving assets for trans- and post-attack command and control.
2. Packet-switching, end-to-end network security, and distributed knowledge and data bases are key technologies needed.
3. A capability is needed for reconstituting remaining partitioned assets, resources and data bases. An automated, possibly knowledge-based, network management capability is needed. Unified internet working protocols and gateways will be required.
4. A layered architecture for organizing the command and control systems applications functions is possible and could aid in developing needed applications and systems software. An extension of an OSI approach to layering appears useful.

6.3 Tactical Command and Control-Interoperability and Survivability

1. For the 1990's, tactical C² systems will need to be interoperable across the military services and their many operational and support systems.
2. Principal requirements are robustness, survivability, high degree of interconnectability, high data transmission rates and interoperability.
3. A common-user, tactical information exchange structure is needed.
4. A tactical system communications protocol reference model and suite(s) of protocols are needed. This should be a joint services effort.
5. A need exists for development of network management and control concepts and protocols, especially with respect to precedence, security and auditing features.
6. New concepts are needed, like the mobile Cellular Command Post, to ensure the survivability of a command center in a tactical nuclear environment.

7. Issues are raised as to how to distribute and manage information data bases and processing resources in a real-time environment.
8. A reconstitution problem exists, similar to that identified for strategic C3. Automatic reconstitution capability is a need to support this.
9. Computer processing machine resources will exhibit a high degree of heterogeneity. A need exists to solve this interoperability problem.

6.4

Protocol Standards for Military Use

1. No single technology is ideal for all applications, yet the full collection of systems must interoperate.
2. DODI 4120.20 requires the adoption of industry standards as DOD standards, in lieu of development and promulgation of new documents. Special exceptions can be made.
3. Currently, use of DOD's TCP and IP protocols appears justified to ensure interoperability when interworking through wide area networks. Intra-LAN use needs further consideration.
4. A new Multinet Gateway approach is intended to normalize the use of diverse wide area networks with the TCP/IP Internet protocols.
5. For the longer term, DOD should carefully consider migration to interoperate with OSI-based systems, as these new protocol-based machines begin to predominate worldwide. This is particularly important for the higher layer distributed processing protocols.
6. The Air Force, in 1983, formulated a policy stating the TCP and IP protocols as being its standards for connection-based transport and internet services within packet-oriented LAN's. In furtherance of this, the Air Force was developing a LAN architecture and established the AFLANSPO Joint LAN Program Office to lead this activity.
7. Air Force priority for use of TCP/IP appears to be attainment of interoperability at some cost in performance. Intra-LAN use of TCP/IP needs further consideration.

6.5

Protocol Standards for Industry Use

1. The International Standards Organization work, in developing the OSI/RM and its suite of protocols, was seen as the key indicator of where industry was going with protocols.

2. Secondly, the direction IBM is taking in extensions to its SNA and use of new LAN technology was seen as another important indicator.
3. The OSI/RM is intended for interoperability among heterogeneous end systems or machines, whereas the IBM/SNA is intended for homogeneous systems and products.
4. Since the DOD C³ environment was shown to be highly heterogeneous, the OSI/RM work and its international scope seem highly relevant to DOD C³ considerations.
5. In spite of OSI, it is expected that in response to users' requirements, the systems built on heterogeneous architecture will grow in number and size. These systems, however, will be open, in that by use of OSI protocols they will be capable of cooperating with other systems. OSI will become the interoperability systems gateway between heterogeneous systems.
6. IBM's SNA, by way of its new Logical Unit 6.2, is now taking on the structure of a truly distributed operating system (DOS). This will influence the architectures of other DOS systems to be developed and is worth tracking. SNA's packaging of protocol layers with a Logical Unit subsystem has merit.

6.6

Tactical Air Control Command, Control and Communications

1. Emphasis on concepts such as behind-line-of-sight target acquisition and weapon delivery techniques is generating an increasing requirement for fast, accurate exchange of information across both functional and service boundaries.
2. A shared use, interconnected, multiple link capability is needed to provide efficient, survivable information transfer services for future C³I users.
3. In TACS, the overall operational management is centralized, but control of the execution is decentralized. A robust/survivable and secure information exchange capability is required.
4. The Air Force's Master Plan for the 1990's (TAFIIS) calls for a dispersed, distributed, survivable C³ system. The plan's concept calls for configuration in a distributed, modular architecture with sharing of information seen as the key to surveillance and intelligence effectiveness.

5. Lightweight, flexible, distributed Modular Operation Centers have been recommended to use LAN's for interconnection to achieve survivability. The use of multimedia to interconnect these LAN's will increase survivability and interoperability.

6.7

C³ for the Tactical Army System

1. A need exists to provide a high degree of survivability and continuity of operations for execution of the air/land battle through use of a distributed tactical C³ system.
2. Distribution and dispersal of command posts are required along with their data bases.
3. The Command Post Network would be a LAN-based system and exhibits characteristics like that discussed for the TACS C³. Common technologies are needed for both missions' systems.

6.8

Distributed Information Processing for C³

1. Command and control applications and systems functionality should be structured with a layered architecture.
2. Heterogeneous processors need to be interoperable over networks, employing high level protocols and packet-switching.
3. A distributed internetwork operating system is required, along with a set of generic C² software.
4. Automated systems resource and network management is needed.
5. C² applications will include real-time, delay-tolerable, delay-insensitive, and delay-variable sensitive types.
6. Digital data, voice, video, imagery, and nondigital forms of information/data must be handled.
7. The architecture must support the interconnection of dispersed clusters or cells of LAN-based processors into a global system.
8. Three levels of operating systems are required: constituent, cluster and intercluster (global).
9. Management of the distributed processes and resources is seen as a major need.
10. Protocols are needed to create a virtual network on top of the distributed, basic physical network, for file access and transfer, data base access and management, resource management, system resource monitoring, system fault-tolerance and survivability, interactive user access and others.

11. Enslow's criteria for a Fully Distributed Processing System (FDPS) should be the basis for the systems architecture:

- a. Multiplicity of general purpose resource components
- b. Physical resource distribution through interconnecting networking
- c. High level operating system
- d. System transparency
- e. Cooperative autonomy

6.9 Operating Systems for C³ Distributed Information Processing

1. A global NOS or DOS form of high level operating system is required to manage the distributed system resources. The NOS form builds on top of the original heterogeneous constituent or Local Operating System while the DOS replaces it with a uniform homogeneous global one.
2. The National Software Works and the Cronus are examples of the NOS form. IBM's SNA, with LU 6.2 services, appears to be evolving to the DOS form.
3. Object-based system architecture is a new concept in structuring of operating systems.

6.10 National Software Works Network Operating Systems

1. The NSW is one example of the NOS form of global, or high level, operating system, implemented on a WAN (ARPANET).
2. It provided users a uniform access to network-provided utility services for accessing objects such as data, files, programs and computing services distributed around the network on heterogeneous machines.
3. The NSW interprocess communications exhibited the following characteristics:
 - a. Front-end to/from Work Manager
 - short, infrequent, among unrelated processes
 - b. Tool/Foreman to/from Work Manager
 - short, infrequent, among unrelated processes
 - c. Front End to/from Tool/Foreman
 - more frequent, short, continuing, among related processes
 - d. File Package to/from File Package
 - infrequent, very long, among related processes

6.11

Cronus DOS

1. Cronus is another NOS form of high level operating system but is implemented primarily on a LAN plus internet capability through a WAN (ARPANET).
2. It will provide services for system access, object management, process management, authentication, access control, security (limited), symbolic naming, interprocess communications and system monitoring. Cronus is an object-based architecture.
3. Objects have capabilities which are defined in access control lists maintained by the object managers.
4. A suite of high level protocols is required to enable the distributed resources and the object managers to cooperate and perform services.
5. A lower set of protocols supports the resources and object managers by performing interprocess communications.
6. The Cronus DOS design employs several protocols. Some, like the underlying Ethernet and Transmission Control and Internet Protocols (TCP/IP), are industry and DOD standard forms. These provide the basic interprocess communications. Others, like the MSF, SER, MSL, OS and OOP, are new ones developed for Cronus. These perform higher layer type protocol functions. They correspond to the OSI/RM layers 5, 6 and 7 protocols.

6.12

Generic Network Operating System (GNOS)

1. To avoid the development of a closed global operating system set of networking protocols for C³I, a general representation was developed for a Generic Network Operating System, called GNOS.
2. A GNOS architecture reference model should be the basis for guiding development of C³I protocols which are open in their interoperability.
3. GNOS is described in terms of its architecture, services, functions, subsystems and protocols to form the basis as a reference model for C³ distributed processing.

4. Protocols needed to support GNOS are identified. These are message-based transaction and file transfer stream types. A networking suite is comprised of three service regions:
 - a. Networking-wide Utilities
 - b. Host to Host/Internetworking
 - c. Local/Wide Area NetworkingIn addition, resource manager, process and object protocols are required.
5. The combined resource manage and networking utility protocols, plus the underlying interprocess communications protocols, form the GNOS.

6.13

Comparison of DOD and ISO Networking Protocol Reference Models

1. Networking reference models provide very important architectural definitions for the services, functions, and protocols necessary to accomplish interoperability.
2. The DOD and ISO communities have found it necessary to develop their own respective models.
3. For interoperability, adherence to particular protocols is essential, whereas adherences to the architecture of the reference model will not ensure interoperability by itself.
4. There are advantages and disadvantages to using a layered approach to defining a networking architecture; however, the advantages are considered great, while the disadvantages are slight.
5. DOD and ISO models are similar. The bottom most network/transport services and protocols are quite close, but the uppermost virtual utility type services and protocols have been architected differently.
6. In DOD's upper layers, the functions and protocols are packaged more into vertical subsystems, whereas in ISO's upper layers, they have been more formally structured horizontally into discrete layers. Elements of both are desired though.
7. Some performance improvement can be gained by a soft layering of protocols, through the sharing of information across layer interfaces. IBM's structuring element called the Logical Unit is a good example of packaging otherwise independent formal layer

protocols together into a more tightly coupled task-oriented subsystem.

8. The DOD model provides the more basic networking utility services/protocols, whereas the ISO set is more comprehensive and is directed more formally to apply object-based design for future distributed processing applications.
9. The DOU model does not contain network or resource management protocols, whereas the ISO model is developing a comprehensive set across the full suite of protocol layers.
10. The OSI/RM has now been approved by both the ISO and CCITT for international adoption and represents a far larger base of users than does the DOD/RM. Longer term, this can place increasing pressure on the DOD to support the OSI/RM protocols as a cost-effective means to maintain interoperability.
11. The IEEE 802 Project has developed the currently leading industry protocols for LAN's. These are undergoing review by the ISO for international adoption. These protocols are usable with both the ISO and DOD models as subnets.
12. Three different methods for media access to a LAN and one common link layer method have been developed. Newer work is adding management and internet working capabilities.
13. Currently, the IEEE 802 work has focused on the 10-Mb/s operating speed range for LAN's.
14. Newer work in the IEEE 802 is looking at larger sized networks than LAN's called Metropolitan Area Networks, or MAN's.

6.14

ANSI 100-Mb/s LAN

1. In the 100-Mb/s speed range for LAN's, the ANSI X3T9.5 Committee is developing the Fiber Distributed Data Interface (FDDI) standard. This is to be a Token Ring LAN using fiber-optic cable and will be based on the IEEE 802 Token Ring.
2. The FDDI will have application in not only front-end LAN's, but back-end high intensity computer/storage data transfers.
3. The FDDI could have application potential for military C³ functions as an alternative to the FILAN.
4. The FDDI exhibits a set of open protocols.

Air Force Flexible Intraconnect LAN (FILAN)

1. The FILAN, currently in an Advanced Development state, is a 180-Mb/s bus-oriented LAN intended for high quantity use in C3I applications. It provides speed capability at a range beyond the IEEE 802 (18 times) and ANSI FDDI (1.6 times) LAN's.
2. The FILAN services span the same layers as the IEEE 802 and goes beyond the FDDI to a degree. There is extensive System Network Management functionality built into FILAN.
3. Extensive capacity, modularity, flexibility and deterministic control features are built into the FILAN.
4. A MIL-STD-1779 User Interface has been established to formalize the method of interfacing to the FILAN Network Access Units to gain service.
5. The FILAN's developed protocols are considered to represent a closed, rather than open, set. Gateway interface devices would be needed to interoperate with industry standard LAN's, such as IEEE 802 and/or ANSI FDDI.
6. Extensions to FILAN are considering the use of mixed media cable and radio paths.

LAN Protocol Characteristics and Effects

1. There are three principal topologies employed to build LAN's: star, bus and ring. Each offers its own set of attributes and limitations.
2. Of the three, broadband bus-oriented LAN's offer the best overall capability for handling a mixture of voice, data and video (employing modulated carriers on broadband coax cable) but rings offer the opportunity to employ wide bandwidth fiber-optic point-to-point cable. Initially, broadband coax bus LAN's will outnumber fiber optic rings, but in the longer term the reverse is expected to occur.
3. For tactical deployments in C3 applications, a mixture of LAN media is going to be required, not only to exploit the use of coax and fiber optics, but also radio paths as well. A Multimedia LAN study currently is under way for RADC by Harris Corporation looking into these issues for extending the FILAN.

4. Both connection-oriented (virtual circuits) and connectionless (datagrams) services are required of LAN's. The IEEE 802 and ANSI FDDI protocols are being developed to support both of these services. For distributed processing, the transaction mode of inquiry/response using datagrams will be the dominant form of message exchange.
5. Broadcast and multicast delivery service, in addition to the basic singlecast service, are readily supported by bus and ring LAN's.
6. In LAN's, communication bandwidth efficiency is traded off to reduce delay caused primarily from protocol processing functions. This is just the reverse from the case of WAN's. More powerful protocol processing machines will be needed in LAN's, along with the wider bandwidths at the mediums, to support C² application.
7. A systematic performance prediction and assessment capability is needed to enable systems engineering of the expected complex LAN-based configurations. A simulation/modeling tool is one of the methods, that, if employed in this process would increase productivity.
8. A variety of media access methods that exist to control how LAN users share the bandwidth resource made available fall into two categories; contention and deterministic. Both can employ distributed control methods.
9. For light to moderate traffic loads, the contention access method(s) provide a lower cost solution with low delays achieved, such as CSMA/CD (Ethernet form).
10. Where traffic loads are expected to be moderate to high and/or where controlled delay has to be a stringent criteria, some form of deterministic access method has to be employed, such as Token Bus or Token Ring.
11. Based on analysis, the CSMA/CD offers an acceptable contention scheme for light to moderate loads while the Token Access Methods do for the deterministic approach when moderate to heavy loads and controlled delay are criteria. The Token Ring is superior to the Token Bus by a small margin, in general.

Evaluations of TCP/IP in a LAN Environment

1. MITRE evaluation reports, on TCP used in LAN configurations, indicated throughput capabilities of 350 kb/s (4096 bit message size) and 900 kb/s in separate measurements. This formed a basis for predicting a 1 Mb/s throughput rate being achievable with a 10-Mb/s Ethernet LAN.
2. A "discretionary TCP/IP" enhancement approach, considered also by MITRE, indicated a way to use only LAN essential elements of full WAN-based protocols, with performance improvements.
3. A discrete-event simulation model can provide a flexible tool for the evaluation of an integrated protocol suite, layered together in a LAN environment. While not fully completed, the model developed for the study represented adequately the Ethernet Blue Book IP, UDP and most of the TCP functions. Completion of the TCP special functions, DOD higher layer protocols and Token Bus/Ring LAN subnets were deferred, in order to conserve funding resources.
4. The throughput attainable with an Ethernet contention type LAN access method were reduced as much as 50 percent as the loading on the cable varies from light to heavy conditions. In addition, while one-way delays were low and constant for light to medium cable loadings, at around 65 percent loading delay characteristics increased sharply. For maximizing total system throughput at the cable reference point, very large message sizes yielded the best results. The reduction in throughput as loading was increased, and/or message sizes were reduced, was caused by increasing collisions which were produced at the cable level. These effects were passed through upward to be reflected in degraded performance of the TCP/IP higher layer protocols.
5. Throughput and delay performance of TCP and IP were very sensitive to different representations of internal machine CPU processing speed, independent of protocol overhead. Faster CPU speeds yielded correspondingly greater throughput and lower delays. This was independent of the effects of the Ethernet collisions. Throughputs of 0.65 through 5.2 Mb/s were demonstrated for TCP-TCP over a single connection without any collisions, for large message size (11K bits). The buildup of queue size at the TCP receiving station was found to limit performance.

6. In multiple-node multiple-TCP connection configurations, throughput for large and small messages at TCP was maintained constant and unaffected by total system loading at the Ethernet cable at up to 65 percent capacity. Thereafter, throughput decrease occurred and delay increased. Ethernet collisions seemed to be the cause of this.
7. Protocol overhead, produced when control header bits were added to the TCP client, or User message sizes, varied as follows:

<u>TCP Client Message Size</u>	<u>Combined Overhead</u>
8 bits, 1%	99%
512 bits, 46%	54%
10,240 bits, 94.5 %	5.5%

6.18

TCP Alternatives for Intra-LAN Use

1. Three references reported are alternatives to using the full TCP for intra-LAN operation. The three selected other than full TCP; one chose a subset, while the other two departed completely from TCP.
2. In one approach, a Local Network TCP was selected as a compatible subset of the WAN TCP, based on a philosophy of expending communication channel bandwidth in order to reduce protocol processing time. This approach substitutes for the TCP flow control, resequencing and duplicate detection features, and uses a much larger packet size than in a WAN. Both the LNTCP and TCP would be implemented in each LAN node, sharing software for the same functions. This would provide for efficient local intra-LAN performance while allowing for good inter-LAN (via WAN) performance as well.
3. An alternative to using TCP at all took an entirely new view to intra-LAN communications, by applying an extension to intra-machine backplane communications in its place. This employed a transparent remote procedure-call form of message-based interprocessor communication in support of a distributed operating system. This method was chosen as an alternative to using internetwork protocols on a LAN, which would lead to minimal benefits and significant costs. To satisfy the needs of distributed processing operating system requirements, the procedure-call-like request-response

SECTION 7.0
RECOMMENDATIONS

7.0 RECOMMENDATIONS

This section presents a roadmap plan in the form of a set of recommendations for consideration by the Air Force, based upon the study's conclusions presented in Section 6.0. The recommendations address the following:

1. Joint Air Force-Industry C³I Applications and Protocol Development Effort.
2. C³I Applications and Systems Layered Reference Model.
3. Generic Network Operating System (GNOS) for C³I.
4. Use of gateways and GNOS Protocols for LAN-based System Interoperability.
5. Need for different Intra-LAN and Inter-LAN Protocols for Underlying Transport Services.
6. Multiple Industry Standardized LAN's to meet Application Requirements.
7. Continued research into Protocol Design, Validation and Formal Verification Methods.
8. Design Practices Handbook for Quantifying LAN Performance characteristics.
9. Performance Evaluation of High Level, Multiple LAN's and WAN Protocols through Simulation and Modeling.
10. Future Study in areas of: 1) Distributed System Design, 2) Integrated OSI for Distributed Information Processing and 3) Multilevel Secure Distributed Operating System.

1. An integrated and coordinated joint Air Force-Industry effort to developing applications of distributed processing, operating systems, data base management and networking protocols for C³I Systems is recommended.

One approach for consideration, based upon experience gained from participation in the IEEE 802 LAN and SAE AE-913 High-Speed Data Bus standards bodies, would be to set up a C³I Architecture and Protocols Development body of committees with appropriate working groups. This body would be led by the Air Force but have active participation from industry contractors and researchers. The scope would include policy, requirements, applications, architecture, distributed processing, distributed operating systems, data base, security and networking protocols (spanning all layers from media through applications layers).

Government funding and IR&D resources are sources for supporting participation. Not only would research and technology areas be addressed but also translation into finalized solutions, designs and military standards/

recommendations affecting procurements for C³I programs would be the real output. Coordination with strategic C³I and the other tactical services would be necessary to assure interoperability for C² information sharing and systems survivability.

2. A layered architecture reference model for C³I systems and applications development is recommended.

The complexity associated with C³I systems can benefit from the use of a layered architecture reference model. This model would divide up the major C³I functions into groupings of layers such that end user applications reside at the top and physical resources reside at the bottom most. The overall model would be structured to identify C³I applications, C² utilities, data base and the Generic Network Operating System (GNOS) containing the networking protocol layer regions (application layer through medium). The model would be defined formally in terms of overall C³I architecture, services, functions, subsystems, interfaces and protocols. Such a model would be a basis for guiding the work of the C³I Architecture and Protocols Development body (Recommendation 1 above).

3. A Generic Network Operating System (GNOS) is recommended to be defined for and guide C³I systems developments.

The heart of the recommended C³I Layered Architecture model is the distributed global operating system and its suite of protocols. A proposed GNOS, described in Section 5.0, Results, and Appendix D, would provide an open, vendor independent model and formal definition to guide this part of the C³I architecture. Since real implementations of C³I systems will entail complex interoperability among heterogeneous elements and end systems, GNOS is needed to exhibit the properties of an open system.

An important part of the GNOS will be the mechanism and procedures necessary for multilevel security in the networking environment and the various protocols. Other important aspects are the user interfaces, programming languages, an operating system command and response language, object to object peer protocols, resource manager to resource manager peer protocols, networking utilities protocols, host to host and internetworking protocols and the underlying LAN's and WAN's protocols and gateways. Section 8.0, Areas of Further Study, discusses many of these issues as do the appendices.

4. Standardized use of generic gateway elements and open GNOS protocols are recommended as the fundamental approach to LAN-based systems interoperability.

A family of gateway interconnection elements is needed for joining together the expected mixture of heterogeneous devices, machines, LAN's and end

systems which will exist and be used in building C³I systems. These gateway elements, discussed in Section 5.0, Results, and further defined in Appendix F, span from interconnecting mediums at the cable through entire end systems at the application protocol layer. A family of gateway elements should be defined and standardized for appropriate levels of interconnecting.

Additionally, it will be through adherence to specific peer protocols where interoperability will be obtained. These protocols should be open (nonproprietary) and exhibit virtual canonical end to end services. While the DOD currently has a network protocol reference model, for the longer term the reference model and protocols of the ISO are expected to require supporting, as manufacturers' machines and systems more and more internationally incorporate ISO's OSI protocols. Further, the ISO protocol development work is considered to represent a far more advanced approach than that of the current DOD towards supporting distributed processing and distributed operating system services and protocols along the lines needed for C³I use. A global OSI systems approach is now under way in the restructured work of ISO.

Packaging of suites of the high level protocols should take account of the way IBM LNA's Logical Units (LU'S) are grouped into task or end user oriented subsystems. This packaging of appropriate elements of multiple protocol layers seems directly beneficial to a distributed operating systems architecture with performance improvement resulting from "soft layering".

As the current DOD high level protocols provides the more basic services for application usage, the richer set of protocols and range of services enhancements seen under development in the ISO work should be the basis for C³I networking utility protocols and services. The military C³I suite of protocols should incorporate the ISO's OSI protocols as a supportable subset. Additions to OSI protocols should be the goal for C³I, not total replacement from developing closed protocols different from common OSI services.

5. Protocols for Intra-LAN and Inter-LAN Transport through Subnet Services are recommended to be separate and distinct.

Distributed processing and operating systems constructed on LAN's need protocols, in the medium through transport layers of the network reference model, which have different characteristics from those being employed today to interconnect machines using a WAN. Most traffic will be for intra-LAN and less for inter-LAN exchanges. Relay gateways and bridge gateways can be employed which will eliminate the need for internet protocols in many instances. The predominant

form of interprocess communications to be employed appears to be short, transaction-based remote procedure calls for inquiry/response. This can be provided by a connectionless (or data-gram) service. However, a connection-based service will also be required, where streams of data (a file) is to be transported. This presents a different form of interprocess communications from that exhibited in today's DOD mandatory Transmission Control Protocol (TCP).

In support of this, the IEEE 802 LAN and ANSI X3T95 FDDI protocols are based upon requiring the connectionless service be mandatory with a connection form optional. Additionally, these inherently support single cast, multicast and broadcast forms of intra-LAN communications. Robust connection-oriented transport protocols should be implemented in LAN to WAN gateways and provide the end to end reliability through that portion of our internet. Translation from LAN to WAN should be confined to the gateway characteristics.

Further, in a LAN environment, protocols should be designed to tradeoff bandwidth efficiency in order to reduce protocol processing times. With LAN speeds moving from 10's to 100's of megabits per second, and faster and cheaper microprocessor and VLSI/VHSIC devices, this tradeoff can produce even greater performance in a LAN environment along with a whole new philosophy for structuring protocols for LAN's.

Development of these protocols should be done in close cooperation with the industry standards bodies, such as IEEE 802, SAE/AE-98 and ANSI X3T95.

6. Multiple industry standardized LAN types are recommended to be employed for matching applications requirements.

No single LAN will satisfy the diverse range of C³I requirements. On the other hand, LAN's should be employed where their unique characteristics are superior, overall. Use of vendor-independent, nonproprietary, or open LAN's is suggested. This can be achieved by basing selection upon industry standardized LAN protocols. Currently, the IEEE 802 represents the 10 Mb/s class while the ANSI X3T95 is developing a 100 Mb/s class. Other organizations are known to be starting work in the 100-1000 Mb/s class (SAE-AE9B Avionics High-Speed Data Bus). Where lacking in services or features, the Air Force should add on to industry LAN standards so as to create a compatible superset. Maintaining as much compatibility with the recognized industry standards will be both equipment cost-effective as well as enhance LAN interconnectability and interoperability.

As a good rule of thumb, the following represents suitable uses for the three leading LAN media access methods developed by IEEE 802:

1. CSMA/CD (Ethernet or contention access)

Smaller sized 10 Mb/s LAN's, where traffic loadings don't exceed about 50-60 percent and where low or controlled maximum delays are critical.

2. Token Access (Deterministic Access)

Larger sized 10 Mb/s LAN's, where traffic loadings will exceed 50-60 percent and where low or controlled maximum delays are critical. Token Ring in addition supports multiple priorities and a mix of synchronous and nonsynchronous traffic for mixing voice and data.

3. Broadband Coax

This, with its inherent wide bandwidth and multiplexed channelization, supports voice, data and video simultaneously.

7. Continued progress is needed in the areas of formal protocol design, validation and formal verification methods. It is recommended to sponsor study, demonstration of methods and development of recommended design practices, in conjunction with work already under way in industry.

Protocols, comprising the heart of a distributed processing system, are crucial to its proper operation. The asynchronous and nondeterministic properties associated with remotely intercommunicating finite state machine processor and systems leads to our overly complex design and certification problem. Work has already been started on this by Dr. Carl Sunshine and others in ISO and CCITT for communications networking protocols. This needs to be applied equally to other areas above the network communications protocol layers to take account of resource (object) manager and process to prosiac protocols.

The impact of VHSIC/VLSI in the implementation of protocol-based distributed operating systems needs to be taken into account as well.

8. Quantification of LAN performance characteristics and guideline design practices for LAN's are recommended. This should employ combined analytic, simulation and measurement methods.

Engineering analysis and practical prediction methods, needed for understanding, comparing and designing LAN's, are not readily available to developers and users of LAN's. This causes confusion and unnecessary costs, and can lead to making the wrong choice of a LAN for an application.

Now that the standards work of the IEEE 802 has matured for CSMA/CD, Token Bus, Token Ring Access Methods and the common Logical Link Control, it would be possible and beneficial to develop a set of recommended design, analysis and

prediction practices covering all of these. A combination of analysis simulation and measurement approaches is suggested. This needs to cover the properties affecting throughput, delay, topology sizings, message size variations and other important properties. Comparison and suggested ranges for applications could be made as well. The output of this would produce a designer handbook.

9. Simulation and modeling effort for evaluating performance of high level, multiple LAN and WAN protocols for C³ is recommended.

The Lan Interoperability Study contract was only able to support the completion of a small portion of the original objectives set for the simulation and modeling of LAN-based protocols (see Table 4.8). This work should be continued to enable quantifying the performance of integrated protocols spanning all layers, under the loading effects of a distributed and/or network operating system set of user processes.

The model should add IEEE 802 Token Bus and Token Ring access methods as well as take into account connection-oriented services for the Logical Link Control. Full TCP and IP services need completing. Next, high level protocols for terminals, files, jobs, mail and name sender would be added and evaluated. Alternatives for TCP and IP for Intra-LAN operation should be evaluated. The sensitivity to varying processor speeds on protocol delays should be examined further. The model should then be expanded to add an internet between LAN's and investigate internetworking effects and gateway operations. The findings should be compiled and published to assist developers.

10. Several additional technological areas of study are recommended to be supported.

Section 8.0 identifies and discusses three areas of future work. These are as follows:

- a. Distributed Systems Design
- b. Integrated Open Systems Interconnection for Distributed Information Processing
- c. Multilevel Secure Distributed Operating System

In general, attendance at the 1983 and 1984 Distributed Systems Technology Exchange meetings at RADC highlighted an apparent lack of interstudy coordination. The areas reported on included LAN's, networking protocols, distributed operating systems, data base, knowledge-base systems, multilevel security and multi-media workstation technologies. There was lacking an overall framework, a common systems architecture for tying these technologies together.

As a result, there were different assumptions about this "real problem," a duplication of functions and a grouping of advocates into distinct camps. Recommendations 1-9, if implemented, would solve these problems.

The three areas of future work discussed in Section 8.0 should be pursued with the above considerations in mind.

SECTION 8.0
AREAS OF FUTURE WORK

8.0 AREAS OF FUTURE WORK

8.1 General

This section presents a discussion on three areas where future technological investigative type work is needed. The three are as follows:

1. Distributed System Design Issues
2. Integrated Open Systems Interconnection for Distributed Information Processing
3. Multilevel Secure Distributed Operating System

While there are many other areas needing further study, too, this report focuses upon the three identified ones because of their relative importance.

8.2 Distributed System Design Issues

8.2.1 Introduction

There are a number of unresolved issues surrounding the subject of the design of distributed systems. These range from architectural issues to terminology disputes.

There are at least three types of distributed systems, which can be characterized as follows [26, 27]:

Network Operating System - a system in which a group of two or more host machines are interconnected over a network communications medium; the hosts are loosely bound by the network operating system (NOS) which manages the resources of the distributed system in a cooperative way, i.e., via request rather than via directive; the user may use the resources of the local hosts via interaction with the local constituent operating system (COS), or the resources of the distributed system via interaction with the NOS; the COS's may be all of the same type, e.g., all VAX-11's running VMS, in which case the distributed system is known as a homogeneous NOS, or each COS may be completely different, e.g., a VAX-11 VMS two IBM-PC's, a PDP-11/45 RSX-11M system, and an Apple PC, in which case the distributed system is known as a heterogeneous NOS.

Distributed Operating System - a system in which a group of two or more host machines are interconnected over a network communications medium; the hosts are under the direct control of the distributed operating system (DOS), which is the only operating system (i.e., there are no COS's); the DOS is directive in nature, able to issue explicit directives to the host machines; while the actual hardware resources may be composed of identical or dissimilar devices, the nature of a DOS is homogeneous, since the user interface is the DOS, of which there is, of course, only one.

Distributed Processing System - a system in which a group of two or more host machines are interconnected over a network communications medium; the user processes communicate with each other via the communication network protocols, and coordinate their functions on an internal level; there is, therefore, no distribution of operating system functions or system control, but this does not preclude directed use of distributed system resources.

These definitions are not universally accepted. For instance, the Cronus system [29,30,31] is characterized as a "distributed operating system" in its support documentation, while using the above definitions, it would be classified as a network operating system. Also, the characterization of homogeneous or heterogeneous could be based on the underlying hardware elements, rather than on the local system control software.

It is possible to have some hybrid combination of any of the above definitions. For example, a system might basically support a DOS view for host machines, but have semi-autonomous terminal concentrators to save the DOS the overhead of managing the internal functioning of these devices.

8.2.2 Overall Architectural Model

8.2.2.1 Viewing Resources as Objects

A flexible and powerful approach to defining the architecture of distributed systems is the object oriented approach. In this approach, all of the resources associated with a computer system are represented as specific instances of a limited number of object classes. Each class of objects encompasses a group of resources or entities which share common characteristics. Basic object classes include processes, files, and the various types of devices. Each object has associated with it specific, well defined operations which may be performed upon or through use of the object. Each specific instance of an object, such as a user data file, has a name or other identifying mechanism, and other components which allow it to interact with the outside world. These may include access rights lists, specific action prevention authorization (e.g., read-only access, etc.), performance limitations (e.g., memory allocation limits, etc.), and similar descriptive elements.

Each object is associated with an "object manager" which knows all the details regarding use of the object. The object manager serves as the intermediary between the user of the objects resources and the object itself. The object manager, therefore, can ensure that all of the generic rules of the object

class (e.g., printers can print, but they cannot read), and the rules of the specific instance of the object (e.g., only users X, Y, or Z can write to this file) are enforced. The operating system, at all levels, can be thought of as the complete collection of object managers.

The object oriented approach is particularly useful in the design of distributed systems. The modular nature of the object model provides a flexibility which is invaluable in distributed system implementation.

8.2.2.2 Distributed System Components

A distributed system can be partitioned into a number of components, including:

Organic and Other Processes - all users of the system are processes; any action which the system takes is taken at the direction of some process.

Programming and Data Manipulation Languages - including FORTRAN, Pascal, and other languages

Operating System Command and Response Language (OSCRL) - including interactive command language interpreters, and statements which can be embedded in language commands (e.g., JCL for IBM system, or Executive Directives for DEC VAX/VMS and PDP/RSX-11 systems)

Communications Networks (including OSI) - the underlying communications hardware, software, and protocols used to connect the local computer systems into a network

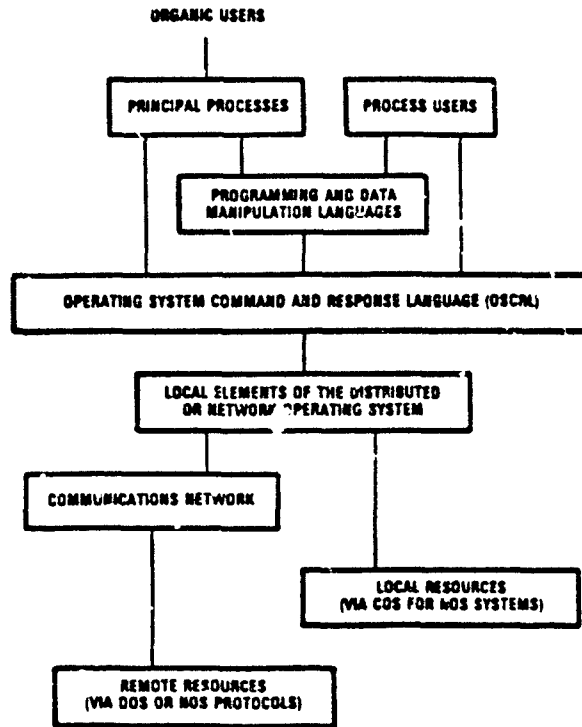
Distributed Resource Control - the remote resource management functions of the NUS or DOS, or the agreed upon principals for remote resource use in distributed processing systems

Local Resource Control - the COS or local resource control portion of a DOS

The organization and function of these components will, of course, vary from system to system. Not every system will have all of the components. For instance, a series of automated data gathering stations operated under the control of a dedicated processor may not have any programming and data manipulation language resources available - it may be implemented in firmware.

8.2.2.3 Relationships Between Components

Figure 8.2.2.3 shows a possible architectural model which can be used to discuss the relationships between the above entities:



13457-46

Figure 8.2.2.3. An Architecture Model for Distributed Processing Systems

The model describes the fact that a computer system exists to allow the users (human and processes) to solve problems using system resources. The intermediate entities provide support for these functions.

8.2.3 Specific Design Issues

8.2.3.1 Applicability of the Types of Distributed Systems

A major distributed system design issue involves selection of the basic architecture as described in the definitions given at the beginning of this discussion. Each of the types of distributed systems has certain advantages and disadvantages which might qualify it for use in one application, while making it unsuitable for another.

The NOS and DOS approaches each have a number of system management layers through which to pass before action is taken with the desired system resource. The performance of a process on an NOS system will most likely be

inferior to that of a process on a single machine system, or a DOS. This is because an NOS will make a request for the desired service of the COS, which will then schedule the request according to its own design. A DOS would have the option of issuing a directive to the target local system. (This scenario could be handled by giving all NOS requests the highest possible priority on each of the local systems.) In a distributed processing system, in which there are no system management layers on the level of an NOS or DOS, process performance could be optimized, with system flexibility sacrificed.

An advantage of an NOS is the capability of independent use of the resources of the local systems in the event of network failure, since each machine will have a fully autonomous COS. It is conceivable that a similar capability could be implemented in a DOS also. A related issue is that by preserving the independence of the local COS's. An NOS could be implemented onto a network of existing systems with a minimum impact, allowing existing software to run without modification.

8.2.3.2 System Security

In a networking system, security is particularly important due to the ability of others to have access to local systems over the network. To ensure an adequate security operation, system and data security must be a fundamental design issue. The purpose of security functions within any information processing system is the protection of the system resources from unauthorized access and use. In some cases, this definition can be expanded to also include limiting or at least identifying any damage done if some unauthorized user does obtain access to system resources.

Access to the system is the cornerstone to any security function. Access control presents a special problem in distributed systems. Since the system is specifically designed to share information to remote locations, this is quite understandable. Physical security of the computer system resources must be maintained at all points of the network. Should any node of the network fall under the control of an unauthorized user, it is possible for that user to attempt to access resources anywhere on the network. Of concern in NOS and distributed processing systems is the possibility that a user may attempt to access resources via the local COS in an attempt to avoid the distributed processing system security functions. Further complicating the NOS and distributed processing system security function is the fact that heterogeneous security policies may be implemented on the nodes of the system. System security issues are discussed in Paragraph 8.4.

8.2.3.3 User Interface and Operating System Command Response Language (OSCRL)

Assuming the use of a common OSCRL over the distributed system, the user would not notice a difference in the interaction complexity between an NOS and a DOS. There would, in fact, be no difference between the implementation of a homogeneous NOS and a DOS, while in a heterogeneous NOS, the OSCRL would have heterogeneous COS interfaces for which to provide translation for. Of course, a user on a distributed processing system may not have a common OSCRL at all, depending on the nature of the system.

From the discussion above, it appears that the OSCRL is a major consideration for a distributed system. Since the OSCRL is the primary user interface to the system, it must be designed specifically to support the desired distributed environment. In general, it is important that the OSCRL syntax not rely upon any location designation to access resources, since distribution of processing activities is a function of the NOS or DOS, the details of which are undoubtedly based in part on dynamic attributes of the system unknown to the user at all times. Of course, the OSCRL might support a location specification capability for instances in which the user desires to define a specific location for an operation over the network. The OSCRL must also be able to support the distributed system security scheme, with provisions made for passing of passwords and keys, and other security related matters.

It is not clear at this time what format an OSCRL for distributed systems should take. Possibilities for the OSCRL range from a command line interpreter, such as DEC DCL, to a sophisticated icon and pointer screen, as in a number of general market microcomputer systems such as Apple Lisa and Macintosh systems. (It is possible that such an icon and pointer user interface could be implemented as a layer above a command oriented OSCRL.) Definite advantages exist for either type of system, again dependent upon the nature of the application.

8.2.3.4 Programming and Data Manipulation Languages

Programming and data manipulation languages are the primary medium through which users direct computer systems. User languages usually interact with the computer system through the OSCRL (which is actually another language). Distributed processing presents a number of challenges to programming and data manipulation language users. Two problems which immediately come to mind are parallel processing and device specification.

If a particular procedure can be partitioned into portions which can run in parallel on different machines within the distributed system, the time required to execute the procedure can be reduced. Unfortunately, the majority of programming languages which exist today do not explicitly support this type of operation, sometimes called concurrent processing. Parallel processing can produce a number of error conditions which must be guarded against. These include the possibility that execution of one of the concurrent modules first may produce different results if this module had not been first. Another potential problem is the possibility that two or more concurrent modules might end up waiting on each other indefinitely, resulting in a deadlock condition. Languages which are designed to support concurrent processing, such as Ada or Concurrent Pascal, have software constructs which are designed to help prevent these occurrences.

Load sharing, a different concept from concurrent processing, can be used in distributed system applications. In a load sharing situation, when one computer system is overloaded, procedures are sent to other nodes in the distributed system for processing.

Another problem incurred in the use of computer languages is that of device specification. Many languages, such as FORTRAN require the specification of the target device in input and output statements. Consider the FORTRAN statement:

```
WRITE (6,10) X
```

The intent of this statement is to write the value of the variable "X" out to device "6" according to the format given in statement number "10". While the variable name "X" and statement number "10" are program specific and independent of the machine, the device specification of "6" is not necessarily the proper device on all systems within the distributed network.

In order for processes to be executable on any computer system within the network, languages must either avoid such specifications, or the distributed system must be able to resolve these references. In some cases, however, reference to a specific device might be desired. It is possible that the distributed system OSCRL could include functions which would allow this issue to be solved on a case-by-case basis.

8.2.3.5 Distribution Issues

Function distribution presents a number of issues which are normally not a consideration in single computer systems. One of these is the degree of

dispersal. Processes within groups of processes, portions of single processes, entire databases, and copies of individual files can be dispersed throughout the distributed environment. Given a five node network, a file which has a degree of dispersal of 100% would have copies resident on all five nodes, while a file of degree 60% would have copies on only three of the nodes. While dispersal of 100% for every process and file in the system would ensure the maximum degree of survivability, this would also be very inefficient. Each modification of a dispersed file requires that all copies of the file be updated. The overhead associated with maintaining a large number of 100% dispersed files over a large distributed system is obviously quite large.

The manner in which nodes are selected for use as dispersal sites is also an issue. Dispersal sites can be selected at random, given a specified degree desired. Alternatively, preset sites for dispersal can be specified for each node in the distributed system. Other possibilities include specification of certain groups of nodes for dispersal of certain types of processes or files, specification of certain nodes to which a given type of process or file is NOT to be dispersed to, and others. An important system service in regard to dispersal is a utility which can be run to modify the dispersal algorithm. For instance, in a battlefield environment, you may want to prevent the dispersal of files to nodes in eminent danger of being overtaken by the opposing force.

Another important issue in distributed systems design is the levels of distribution. How many distributed systems will a given computer system be involved in? It is conceivable that any given system will be involved in distributed systems which involve other computer systems at the same organizational plateau, and others which are not. For example, in the TAC C3 DOS Study [23], a number of processors are networked together under the direction of a distributed operating system known as the Minidos. For upper echelon command levels, a number of Minidos networks are networked together under the cognizance of a Maxidos, which is layered on top of the Minidos. (The characterization of these levels as a DOS is again, not according to the definitions presented at the beginning of this paper. The system described in the TAC C3 DOS Study would, according to these definitions, be characterized as an NOS, since each computer system has its own COS.) While the TAC study stops here, it is not unreasonable to envision a need arising for additional or overlapping layers beyond the Minidos and Maxidos networks.

8.2.4 Conclusions

Distributed system control is a very complex subject which will require a great deal of further research to become fully effective. The inherent flexibility of distributed systems will allow a great diversity in the manner of control of distributed systems. Just as the specific type of distributed system will vary from application to application, so too will the appropriate form of distributed system control.

There are a number of issues and general areas in this field which warrant further research. Several of these are:

- Potential uses of artificial intelligence (AI) techniques for network control and security functions
- The nature and design of an appropriate user interface, both user interactive and process embedded
- Philosophies and algorithms for degree of dispersion and similar survivability and object update and control issues
- Design philosophies and issues involved in multiple levels of distributed control (e.g., Minidos and Maxidos)
- System reliability and performance tradeoff studies
- Distributed system security functions (in addition to those studied under the AI subject area)
- Centralization and decentralization of distributed system control
- System failure and degradation control, both intentional and unintentional
- Distributed system topology and growth management issues

This list is by no means all inclusive. Many of these studies will have to be carried out in light of specific objectives, since conclusions reached for one type of application may not be suitable for others. For example, a system which required quick responses of the user in order to prevent unauthorized and untrained users to access the system would not be suitable for a university teaching environment. As with the evolution of COS's, there will be a multiplicity of distributed systems, some general purpose, others more specifically tailored to a given application.

8.3 Integrated Open Systems Interconnection for Distributed Information Processing

8.3.1 ANSI/OSI Networking Committee Reorganization

Several factors in the field of computer networking and distributed processing are driving a reorganization of the ISO and ANSI committees which deal with these and related technologies. Currently, the areas of computer networking, database technology, data security, and operating system technology are considered distinct and separate fields, each represented by its own committee or subcommittee in the international and national standards bodies.

There is, however, a growing consensus that with the advent of distributed processing technology, the fields are no longer distinctly separate, and therefore should not be separated in standards committee deliberations. The trend is perhaps best summarized by Bachman and Ross [16]:

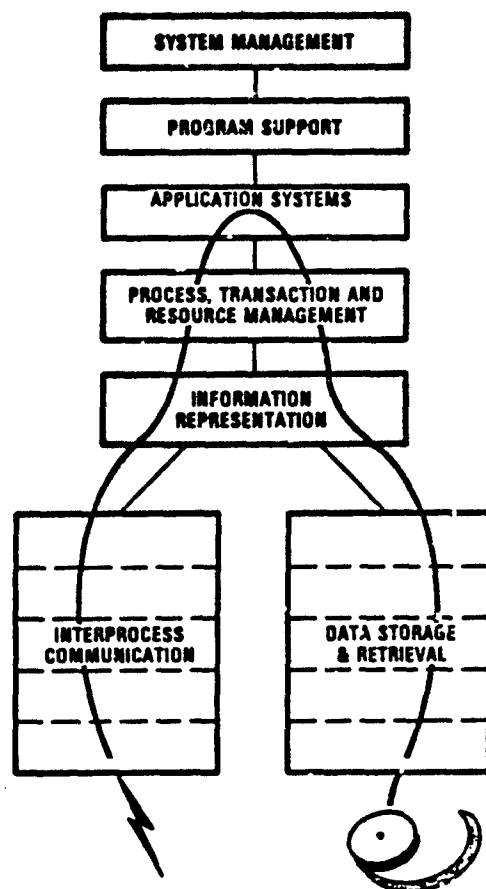
"The ISO Reference Model of Open Systems Interconnection is a major tool of TC97 for the study and organizations of standards activities relating to interprocess communications. The authors believe that the development of a larger reference model to cover the complete scope of computer-based information systems would provide an even greater force within TC97 to assist it in determining the standards to be produced and help place existing standards into a larger context."

The ISO OSI Reference Model is the computer networking communications standard basis currently adopted by the international community. These authors, and others [15], are supporting the idea that more useful standards work could be accomplished by considering this aspect of information processing as only a small part of a larger scheme. Figure 8.3.1 illustrates the enlarged model which reference 16 developed to focus the global OSI work upon.

These changes will therefore bring the full weight of the international standards making organization to the problem of fully distributed information processing technology. It is conceivable that within five to seven years, there may be fully adopted international standards to cover the entire range of data manipulation, operating system and command language interface, and network communications needed to support open system distributed processing.

8.3.1.1 ISO Technical Committee 97 Reorganization

ONNICOM, in its July 1984 Open Systems Communications Newsletter [97], reported on the reorganization of Technical Committee 97 along the lines discussed above, in the manner which follows.



13457-50

Figure 8.3.1. Global OSI Reference Model of Computer Based Information Systems Functional Subarchitectures

"The International Organization for Standardization (ISO) officially restructured its Technical Committee TC97 on Information Processing Systems by an action taken at the TC97 Plenary Meeting in Stockholm, May 16-18, 1984. The restructuring had two principal purposes:

- To improve the manageability of the very large and growing overall program of work of TC97 and the relationship between TC97 and other ISO committees and international organizations (i.e., CCITT and IEC)
- To improve the coordination and liaison between major elements within TC97, in particular the standards related to the upper layers of OSI

In addition to the restructuring, TC97 approved a revised wording of the definition of its scope to increase clarity:

Standardization, including terminology, in the area of information processing systems, including but not limited to personal computers and office equipment.

To improve the manageability and planning of its work program, TC97 was restructured into three groupings, and a Vice Chair was appointed for each grouping:

Application Elements: Mr. A. Tateishi, Canada
Equipment and Media: Prof. H. Wada, Japan
Systems: Mr. Y. Le Roux, France

The Vice Chairs are responsible primarily for ensuring the coordination of, planning for, and interaction among the subcommittees within their grouping. This will promote better harmonization of the standards developed within TC97 and will facilitate liaison with other organizations.

The subcommittees assigned to each grouping are:

Application Elements:

SC 1 Vocabulary

SC 7 Design and Documentation of Computer Based Information Systems

SC14 Representation of Data Elements

Equipment and Media:

SC10 Magnetic Disks

SC11 Flexible Magnetic Media for Digital Data Interchange

SC13 Interconnection of Equipment

SC15 Labeling and File Structure

SC17 Identification and Credit Cards

SC19 Office Equipment and Supplies

SC23 (new) Optical Digital Data Disks

Systems:

SC 2 Character Sets and Information Coding

SC 6 Telecommunications and Information Exchange between Systems

SC18 Text and Office Systems

SC20 Data Cryptographic Techniques

SC21 (new) Information Retrieval, Transfer and Management for Open
Systems Interconnection

SC22 (new) Application Systems Environments and Programming Languages

The changes of greatest interest to OMNICON subscribers are within the Systems grouping. These include changes to SC6 and SC18, and the establishment of the new SC21:

SC6, Telecommunications and Information Exchange Between Systems

Transport Layer is added to the scope of SC6, which is now responsible for the lower four layers of OSI: Physical, Data Link, Network, and Transport. This makes for a very clean interface between the work of SC6 and that of SC21, which has Session Layer and above: SC6 provides all end-to-end "bit pipes", and SC21 provides for the standardized use of the bit pipes to transfer information across the OSI environment.

SC18, Text and Office Systems

Text Preparation and Interchange Equipment, and Computer Language for the Processing of Text, are added to the scope of SC18, which is now responsible not only for the functional and structural standards for text preparation and interchange (i.e., terminology, document architecture, text processing functions, text layout, and document interchange), but also for specialized terminal equipment standards and programming languages required in this area. This is consistent with the increased application emphasis of the TC97 restructuring, intended to provide better responsiveness of the ISO TC97 standards to the application areas for which the standards are being developed.

SC21, Information Retrieval, Transfer and Management for Open Systems Interconnection

A major new subcommittee has been formed out of major sections of the old SC5, Programming Languages, and SC16, Open Systems Interconnection. SC21 is responsible for the Session Layer and above of OSI, together with computer graphics, database, and operating systems command and response languages. This puts into one subcommittee all standards related to "virtual information resources:" terminal functions, graphics, database, files, OSI, and operating systems services. This should make it far easier to develop and maintain compatibility among all these "upper layer standards."

8.3.2 Integrated OSI for Distributed Information Processing

Combining Communications with Programming Languages, Data Base and Operating System Command Response Language (OSCRL) projects will require work on Open Systems Interconnection for total information processing (Figure 8.3.1) in the context of a larger reference model. The new enlarged model would cover the complete scope of distributed computer-based information systems integrating the layered architectures for data storage/retrieval as well as for interprocess communications and systems management. A common information presentations layer would provide syntax representations for: 1) interprocess communications, 2) data storage and retrieval, and 3) operations on the data local to the process.

The global OSI/RM view in Figure 8.3.1 identifies a new layer for Application Management functions, that are not application type specific, which deal with process integrity and security, called Process, Transaction and Resource Management (P, T and RM Layer). Above the P, T and RM layer, would be the Applications Specific layer entities, while above them would be Software Program Support and Systems Management. Beneath the P, T and RM layer, would be the interprocess communications and the data storage and retrieval subsystems of layered protocols.

High level protocols for networks may become defined in terms of data types where applications protocols would be defined in terms of the universe of discourse appropriate to the distributed application; if so, then programming languages having strong data typing will be important.

The Air Force needs to participate in the above discussed projects in order to benefit more effectively from the resultant protocols and other standards.

8.3.3 Issues Associated with an Integrated OSI for Distributed Information Processing

1. A common universe of discourse is required between cooperating application processes (same semantic interpretation of information exchanged).
2. Commitment and recovery protocols are required in application layer.
3. Distributed open systems need a common (standard) form of an "Operating System Command and Response Language (OSCRL)" meeting the common requirements of all operating systems.
4. OSCRL needs to adopt the characteristics of a real-time, parallel programming language to function in a distributed environment; its data objects become abstract objects specifying characteristics of information processing resources:
 - Processing time, work space, data storage, input/output devices, communication channels and their costs.
5. OSCRL initiated activities in OSI will be subject to both processing delay and communications delay and necessitate the ability to handle not only these delays but recoveries from failures.
6. OSCRL should satisfy both machine as well as human user needs.
7. OSCRL needs to be concerned not only with job processing but with file and data handling resources as well.

8. OSCRL will also be concerned with the management aspects associated with coordination of Application-Process-Group operations in conjunction with OSI management protocols/functions.

8.4 Multilevel Secure Distributed Operating System

8.4.1 Introduction

A multilevel-secure DOS must meet stringent requirements defined by the DoD Computer Security Center; specifically those of the "A1 class." Class A1 demands a number of functional capabilities in security policy and accountability. Its unique feature, however, is the degree of assurance that must be provided, much of it through a strict formal and mathematical methodology.

Fundamental security requirements for a distributed system - authentication, secrecy, integrity and necessity - are no different from those for any other secure system environment. However, certain aspects of a distributed system present special problems, notably: distribution of resources, sharing of resources, distribution of control, heterogeneous host processors, and interaction between possibly different policy domains.

There are two commonly discussed types of system that fit the Enslow [26] definition of a "fully distributed processing system":

- Systems in which hosts are essentially identical and run a homogeneous set of distributed operating system (DOS) software to provide all functions and services.
- Systems in which heterogeneous constituent operating systems (COS) are interconnected by a global DOS (sometimes called a "network operating system" or NOS) that provides, despite the heterogeneity, a uniform global interface, global object access, symbolic naming, and other components of system-wide "transparency".

Current implemented security policy criteria do not address the above issues. A first result, then, is that either clarification of existing criteria or generation of new criteria is needed. Each specific function or service sought for a secure DOS must be examined in the light of security rules and objectives and either allowed, disallowed, or changed. In the process, some of the security criteria may change as well.

Given a policy that is applicable and sufficient to the DOS environment, we still face a number of design and implementation issues. Some examples follow.

- How are distributed resources and their control partitioned?
- How are labels maintained and secrecy requirements enforced across heterogeneous hosts?

- How is message delivery assured?
- How are distributed data and processes synchronized?
- How are resource deadlocks detected and resolved?
- How can we be sure that our design will support growth in applications, users, and equipment?

Additionally, the impact of solutions to these issues needs to be considered in full view of DOS performance requirements.

8.4.2 CRONUS DOS Baseline

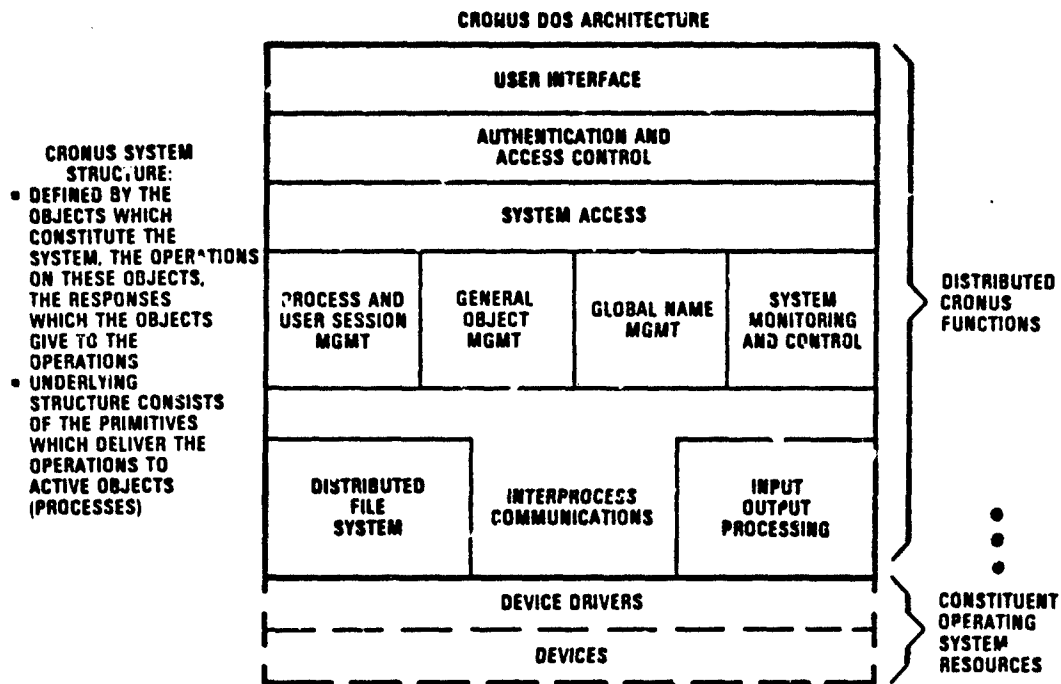
A guiding principle behind CRONUS [29, 30, 31] is computer system interoperability within a LAN-based network of heterogeneous host computers, operating systems, and devices. In CRONUS terminology, this network is called a cluster built on top of a local area network. Despite the wide mix of machines and constituent systems, the user is to see a single coherent DOS providing a range of transparent global services. These include: global system object definition, access control and protection of data and resources, interprocess communication, distributed file management, a uniform mail service, and a network virtual terminal capability. Figure 8.4.2-1 illustrates the single user view of the CRONUS DOS architecture. The multiuser view is given in Figure 8.4.2-2.

The CRONUS cluster enables autonomous heterogeneous machines to perform operations under distributed control as though they were in fact a single large machine. The global reference "window," the Common Command Language (CCL), does not, however, surrender the host machine's independence or autonomy.

In addition to providing transparent interoperability (coherence and uniformity), CRONUS provides benefits directed at other specific design goals: system integrity and survivability, assuring operational continuity despite component failures; scalability in addressing and configuration; global resource management by task priority; component substitution capability; and support for simplified system operation and maintenance.

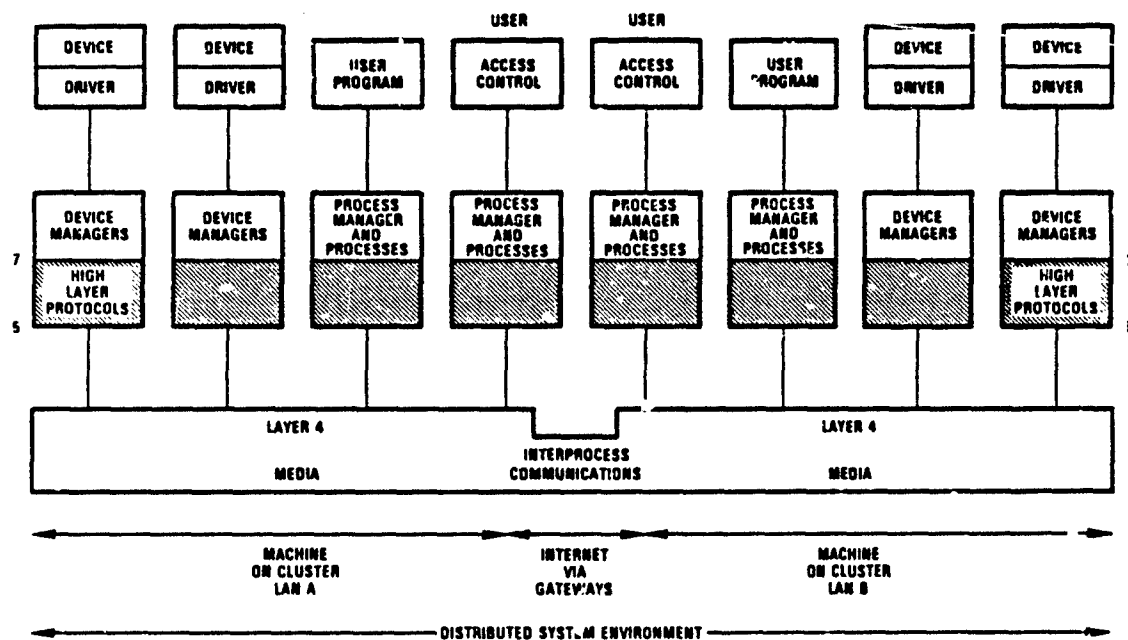
While the CRONUS baseline provides concepts that are readily carried over to a secure environment, the inherent flexibility of the system may introduce potential theoretical and practical problems. Indeed, certain CRONUS elements may be fundamentally unsuited to a secure system design.

The CRONUS user process, system processes, device and other resource managers are interconnected by means of an interprocess (or inter-object) communications network. From a logical perspective, this communication is by way



13422-2

Figure 8.4.2-1. Single User CRONUS DOS Architecture



13422-1

Figure 8.4.2-2. Multiuser Distributed CRONUS System Environment

of peer-to-peer protocols. Procedurally, interprocess messages are composed with the idea of the message structure library and sent through the operation switch. Note that the constituent operating system (COS) is responsible for the direct control of local resources (device drivers, etc.). It also shares aspects of operation of a number of the distributed functions.

CRONUS manages the global resources of the system through an object-based model in which "object" is an abstraction of the resource, to be acted on by a well-defined set of operations. CRONUS objects have unique names that the user refers to without knowing or needing to know either the physical resource involved or the location of the manager for the object. All system interactions can be described by well-defined object operations.

From a security standpoint, a design based on abstract object and on a CRONUS or CRONUS-like object management model can not only provide efficient global handling of resources, but can also strengthen resource isolation and can support an almost unlimited richness of object access controls. And in fact, each instance of a CRONUS object has, with a unique identifier (UID), an object descriptor that defines properties including access rights.

Among the more troublesome areas in attempting to define a secure DOS are the following:

- The very fact of direct terminal user access to a host constituent operating system can, if not controlled in a specific, secure way, compromise CRONUS DOS facilities that rely on the host.
- CRONUS access control lists, a part of the object descriptor and serving to prevent unauthorized use of objects and resources on the system, are not themselves secure.
- User identification of objects and the permanent user data base constructs suffer from a similar vulnerability.
- The operation switch supports migratory and replicated objects, such as files and processes; these objects are moved by the CRONUS system from host to host in an insecure fashion.
- Concurrency control is implemented in such a way that multiple copies of identically named data may be present in the system at the same time. This can lead to evident security compromise.

These problem elements can perhaps all be redesigned for a secure distributed operating system without losing their essential functional characteristics (from a user viewpoint). However, it may be impractical, if not impossible, to secure these elements using the band-aid (patched) approach of modifying the actual CRONUS implementation.

An important principle in computer-based security is that the central protecting mechanism (or complex of mechanisms) must reflect a unified design; it cannot be "patched" into an existing implementation. Nonetheless, existing or proposed systems have capabilities that would or should be preserved in a secure implementation. Among these are:

- Uniformity of procedures for using local and remote resources
- Transparency of heterogeneous host eccentricities
- User access capabilities and limitations
- Enhanced reliability vis-a-vis single host operation

It is desirable to retain as much of this functionality as possible. A secure DOS is not, however, a patched up insecure DOS - it cannot be and still accomplish the unified protection required by the AI-system rules. But the capabilities can be mapped clearly from the existing design to a conceptual overview of what a secure DOS must look like.

Further work needs to be undertaken which examines the above minimally identified issues.

SECTION 9.0

REFERENCES

9.0 REFERENCES

1. Postel, J., "Internet Protocol Transition Workbook," SRI, Int.-Network Information Center, March 1982.
2. _____, "Defense Data Network Program Plan," Defense Communications Agency, January 1982, Revised May 1982.
3. Schneidewind, N. F., "Interconnecting Local Network to Long-Distance Networks," Computer, pp. 15-24, September 1983.
4. Warner, C., "Connecting Local Networks to Long Haul Networks: Issues in Protocol Design," 5th Conference on Local Computer Networks, pp. 71-76, October 6-7, 1980.
5. _____, "International Standard ISO/IS 7493: Information Processing Systems - Open Systems Interconnection - Basic Reference Model," International Standards Organization, 1982.
6. De Lauer, K., "DOD Policy on Standardization of Host-to-Host Protocols for Data Communications Networks," memorandum dated 20 March 1982, Under Secretary of Defense for Research and Engineering.
7. _____, "HQ USAF Local Area Network workshop," 25-27 January 1983.
8. Holmgren, S. F., "Evaluation of TCP/IP in a Local Network," The MITRE Corp., Working Paper 81 W00568, September 30, 1981.
9. Skelton, A., et al., "FY 80 Final Report: Cable Bus Applications in Command Centers," December 1980, The MITRE Corp., MTR-80W 00319.
10. Cerf, V. and Cain, E., "The DOD Internet Architecture Model," pp. 307-318, Computer Networks, Vol. 7, No. 5, October 1983.
11. Selvaggi, P., "The Department of Defense Data Protocol Standardization Program," pp. 319-328, Computer Networks, Vol. 7, No. 5, October 1983.
12. Hanlon, R. "HQ USAF Local Area Network Workshop," 25-27 January 1983.
13. _____, "USAF Local Area Network (LAN) Architecture," Draft, 20 July 1983, HQ USAF/SITT.
14. Folts, H. and Des Jardins, R., "Proceedings of the IEEE - Special Issue on Open Systems Interconnection (OSI) - Standard Architecture and Protocols," December 1983, Volume 71, No. 12, pp. 1329-1486.
15. Enslow, P. H., "Computer Networks - Special Issue on: Programming Languages and Open Systems Interconnection," Volume 8, No. 1, February 1984, pp. 1-55.
16. Bachman, C. W. and Ross, R. G., "Towards a More Complete Reference Model of Computer Based Information Systems," Computer Networks - Volume 6, No. 5, October 1982.

17. Rauch-Hindin, W., "Communication Standards - ISO Poised to Make Its Mark," Systems and Software, March 1984, pp. 104-126.
18. Rosenberg, R., "Closing In On Open Systems," Electronics, May 31, 1984, pp. 78-83.
19. Gray, J. P., et al., "Advanced Program-to-Program Communication in SNA," IBM Systems Journal, Vol. 22, No. 4, 1983, pp. 298-318.
20. Francois, P. and Potocki, A., "Some Methods for Providing OSI Transport in SNA," IBM Journal on Research and Development, Vol. 27, No. 5, September 1983, pp. 452-463.
21. Strole, N. C., "A Local Communications Network Based on Interconnected Token-Access Rings: A Tutorial," IBM Journal on Research and Management, Vol. 27, No. 5, September 1983, pp. 481-496.
22. Pfister, G., et al., "Modular Operations Center Concept Study," AD B066846, ITT Gilfillan, April 1982.
23. Thompson, J. R., et al., "TAC C³ Distributed Operating System Study," AD B046175, Operating Systems, Inc., January 1980.
24. Fearey, J., "System Architectural Concepts: Army Battlefield Command and Control Information Utility (CCIU)," July 15, 1982, Jet Propulsion Lab, AD A124379.
25. _____, "Tactical Information Exchange (TIE) Framework Options Development," Joint Services/OSD/Industry Working Group, October 1981.
26. Enslow, Jr., P. H., "What Is A Distributed Data Processing System?," Computer, January 1978, pp. 3-11.
27. Tanenbaum, A. S., "Computer Networks," Prentice Hall, 1981, pp. 476-483.
28. Schantz, R. E. and Thomas, R. H., "A Technical Overview of the National Software Works Project," AD A132320 and RADC-TR-83-80, Bolt Beranek and Newman, Inc., March 1983.
29. Schantz, R., et al., "Cronus, A Distributed Operating System," RADC-TR-83-236, Bolt Beranek and Newman, Inc., November 1983.
30. Schantz, R., et al., "Cronus, A Distributed Operating System," RADC-TR-83-255, Bolt Beranek and Newman, Inc., December 1983.
31. Schantz, R., et al., "Cronus, A Distributed Operating System: Functional Definition and System Concept," RADC-TR-83-254, Bolt Beranek and Newman, Inc., February 1984.
32. _____, "DOD Protocol Reference Model," Draft, System Development Corporation, PSTP 83-2, January 1983.
33. Lilienkamp, J., et al., "DOD Protocol Reference Model," System Development Corp., TM-7172/201/04, 2 December 1983.

34. Blankertz, W. H. and Gombetz, D. A., "NIS Local Area Network Issues," The MITRE Corp., MTR-82-W00123, July 1982.
35. _____, "Transmission Control Protocol Military Standard," MIL-STD-1778, 12 August 1983.
36. _____, "Internet Protocol Military Standard," MIL-STD-1777, 12 August 1983.
37. _____, "File Transfer Protocol Military Standard," Draft, MIL-STD-1780, 19 September 1983.
38. _____, "Telnet Protocol Military Standard," Draft, MIL-STD-1782, 8 August 1983.
39. _____, "Simple Mail Transfer Protocol Military Standard," Draft, MIL-STD-1781, 8 August 1983.
40. Mills, D. L., "Internet Systems and Protocols," Short Course, George Washington University, January 16-20, 1984, Orlando, Florida.
41. Nelson, J., "802: A Progress Report," Datamation, September 1983.
42. Martin, J., "Design and Strategy for Distributed Data Processing."
43. Cooper, G. H., "An Argument for Soft Layering of Protocols," MIT/LCS/TK-300, May 1983.
44. Gien, M. and Zimmerman, H., "Design Principles for Network Interconnection," Sixth Data Communications Symposium, pp. 109-119, November 27-29, 1979.
45. Skelton, A., et al., "FY 80 Final Report: Cable Bus Applications in Command Centers," December 1980, The MITRE Corp., MTR-80W00319.
46. Summers, J., "Implementation and Application of DOD Standard Protocols in Local Area Networks," 28-30 September 1982, Proceedings of Conference on Local Area Military Networks, Griffiss AFB, NY.
47. Summers, J., "Use of Transmission Control Protocols/Internet Protocols (TCP/IP) in Local Area Network," HQ USAF Local Area Network Workshop, 25-27 January 1983.
48. Holmgren, S., "Evaluation of TCP/IP in a Local Network," September 30, 1981, The MITRE Corp., WP81W00568.
49. Warner, C., "Connecting Local Networks to Long Haul Networks: Issues in Protocol Design."
50. Cheriton, D., "Local Networking and Inter-networking in the Y-System," October 3-6, 1983, Eight Data Communications Proceedings.
51. Schneidewind, N., "Interconnection of Local Network to Long-Distance Networks," September 1983, Computer, IEEE.
52. Stuck, B., "Analyzing Congestion in Local Area Networks: IEEE Computer Society Project 802 Local Area Network Standards."

53. Kummerle, K. and Reiser, H., "Local Area Communication Networks - An Overview," Winter, 1982 Volume 1, Number 4, Journal of Telecommunication Networks Konheim.
54. A. G. and Meister, B., "Waiting Lines and Times in a System with Polling," J.ACM, Vol. 21, 1974, pp. 470-490.
55. Lam, S. S., "A Carrier Sense Multiple Access Protocol for Local Networks," Computer Networks, Vol. 4, 1980, pp. 21-32.
56. Tobagi, F. A. and Hunt, V. B., "Performance Analysis of Carrier Sense Multiple Access with Collision Detection," Technical Report 173, Computer Systems Laboratory, Stanford University, Stanford, CA, 1979.
57. Bux, W., "Local-Area Subnetworks: A Performance Comparison," IEEE Trans. Commun., Vol. COM-29, 1981, pp. 1465-1473.
58. Didic, M. and Wolfinger, B., "Simulation of a Local Computer Network Architecture Applying a Unified Modeling System," pp. 75-91, Computer Networks, Vol. 6, No. 2, May 1982.
59. _____, "The Ethernet - A Local Area Network, Data Link Layer and Physical Layer Specification," Digital, Intel and Xerox, Version 1.0, September 30, 1980.
60. Elden, Walter L., "LAN Based Protocol Issues for Distributed End Systems Interoperability," Fiber-Optic Communications/Local Area Networks 1984 Conference, Las Vegas, NV, September 17-21, 1984.
61. Chapin, L., "Connections and Connectionless Data Transmission," pp. 1365-1371, a paper in reference 14.
62. Gien, M. and Zimmerman, H., "Design Principles for Network Interconnections," Sixth Data Communications Symposium, pp. 109-119, November 27-29, 1979.
63. Bonhamore, E. and Estrin, J., "Multilevel Internetworking Gateways: Architecture and Applications," Computer, pp. 27-34, Vol. 16, No. 9, September 1983.
64. Callon, R., "Internetwork Protocol," pp. 1388-1393 in Reference 14.
65. Ware, C., "The OSI Network Layer: Standards to Cope With the Real World," pp. 1384-1387 in Reference 14.
66. Mier, E. E., "Packet Switching and X.25 -- Where to From Here?," Data Communications, pp. 121-138, October 1983.
67. Elden, W., "Gateways for Interconnecting Local Area and Long Haul Networks," International Conference on Local Networks and Distributed Office Systems, pp. 391-406, May 11-13, 1981, London, U.K.
68. Folts, H., "Internetworking of 802 Local Area Networks and X.25 Wide Area Networks," Open Systems Data Transfer, Transmission #10, June 1984, Unmicon, Inc.

69. Hinden, R., et al., "The DARPA Internet: Interconnecting Heterogeneous Computer Networks With Gateways," Computer, pp. 38-48, Vol. 16, No. 9, September 1983.
70. _____, "Multinet Gateway System Specification for Advanced Development Model (Type A)," Spec. No. C00053, Ford Aerospace and Communication Corp., 21 October 1983.
71. Elden, Walter L., "LAN Interoperability Study-Interim Report," Harris Corporation-GISD, Volumes 1 and 2, January 31, 1984.
72. Elden, Walter L., et al., "LAN Interoperability Study Report (Preliminary Results)," Harris Corporation-GISD, June 18, 1984.
73. Frankel, M. S., "Packet Radios Provide Link for Distributed, Survivable C³ in Post-Attack Scenarios," MSN, June 1983, pp. 80-106.
74. Frankel, M. S., "Telecommunications and Processing for Military Command and Control: Meeting User Needs in the Twenty-First Century," IEEE Communications Magazine, July 1984 - Vol. 22, No. 7, pp. 18-25.
75. _____, "Tactical Information Exchange (TIE) Framework Options Development," Joint Services/OSD/Industry Working Group, October 1981.
76. Cerf, V. and Lyons, R., "Military Requirements for Packet-Switched Networks and Their Implications for Protocol Standardization," Computer Network, Vol. 7, No. 5, October 1983, pp. 293-306.
77. Cerf, V. and Cain, E., "The DOD Internet Architecture Model," Computer Network, Vol. 7, No. 5, October 1983, pp. 307-318.
78. Selvaggi, P., "The Department of Defense Data Protocol Standardization Program," Computer Network, Vol. 7, No. 5, October 1983, pp. 319-328.
79. Henriques, V., Letter to Dr. Edith W. Martin - Department of Defense; Computer Business Equipment Manufacturers Association (CBEMA), September 20, 1983.
80. Latham, D., Letter Response to Mr. Vico E. Henriques - (CBEMA; Department of Defense, 27 October 1983.
81. Snow, D., "A Secure Implementation of the OSI Protocol Layering Model," Compusec, Inc., June 1984.
82. _____, "Systems Network Architecture - General Information," Order No. GA27-3102, "Systems Network Architecture - Format and Protocol Reference Manual: Architectural Logic, Order No. SC 30-3112, available through IBM Branch Offices.
83. Sundstrom, R. J., "Program-to-Program Communications - A Growing Trend," Data Communications, pp. 87-92, February 1984.
84. Joseph, G., "An Introduction to Advanced Program-to-Program Communication (APPC)," IBM Corp., Document No. GG 24-1564-0, July 1983.

85. Francois, P. and Potocki, A., "Some Methods for Providing OSI Transport in SNA," IBM Journal of Research and Development, Vol. 27, No. 5, September 1983, pp. 452-463.
86. Strole, N. C., "A Local Communications Network Based on Interconnected Token-Access Rings: A Tutorial," IBM Journal of Research and Development, Vol. 27, No. 5, September 1983, pp. 461-496.
87. Brick, D. B. and Ellersick, F. W., "Future Air Force Tactical Communication," IEEE Transaction on Communication, Vol. COM-28, No. 9, September 1980, pp. 1551-1572.
88. Thompson, T. H., "Tactical Air Force Integrated Information System Master Plan," Signal, pp. 68-73, August 1978.
89. Pfister, G., et al., "Modular Operations Center Concept Study," AD 8066846, ITT Gilfillan, April 1982.
90. Fearey, J., et al., "System Architectural Concepts: Army Battlefield Command and Control Information Utility (CCIU)," Jet Propulsion Lab, AD A124379, July 15, 1982.
91. Enslow, P. H., et al., "Software Support for Fully Distributed/Loosely Coupled Processing Systems," Georgia Institute of Technology, RADC-TR-83-238 (2 Volumes), January 1984.
92. Cypser, R. J., "Communications Architecture for Distributed Systems," Addison Wesley, 1978, pp. 72-75.
93. Barrow, M., et al., "Cronus, A Distributed Operating System - Interim Technical Report No. 3," Bolt Beranek and Newman, Inc., Report No. 5646, May 1984.
94. Joshi, S. and Iyer, V., "New Standards for Local Networks Push Upper Limits for Lightwave Data," Data Communications, July 1984, pp. 127-138.
95. Miller, A., "Progressive Project Document for Local Area Network Simulation," Harris Corporation - GISD, July 1984.
96. _____, "The Ethernet, A Local Area Network," Digital, Intel and Xerox, Version 1.0, September 30, 1980.
97. Folts, H., "ISO/TC97 Gets Reorganized," Open Systems Communication, Transmission No. 26, July 1984, pp. 7-9.
98. Gliger, V. D. and Luckenbaugh, G. L., "Interconnecting Heterogeneous Data Base Management Systems," Computer, January 1984, pp. 33-43.
99. _____, "Remote Data Base Access Protocol - Third Draft Standard ECMA-DB," European Computer Manufacturers Association ECMA/TC22/84/2, January 1984.
100. _____, "Multimedia LAN Proposal," Harris Corporation - GISD, March 1984.

APPENDICES

This provides two groups of appendices to the report. One group consists of working papers which focus on protocol and design issues needing attention for LAN Interoperability. The second group consists of additional working papers which focus on issues and design approaches to internetworking.

Group 1 Appendices - Protocols and Issues

- A. Evaluation of DOD Higher Layer Protocols
- B. Some Improvements to DOD Higher Layer Protocols
- C. Transmission Control Protocol (TCP) Usage for LANs
- D. Protocols for the Generic Network Operating System (GNOS)
- E. Networking and System Resource Management Protocols
- F. Remote Data Base Access Protocol

Group 2 Appendices - Design Approaches to Internetworking

- G. Generic Gateways for LAN Interoperability
- H. Multi-Media LAN (MMLAN) Internetworking

APPENDIX A
EVALUATION OF DOD HIGHER LAYER PROTOCOLS

APPENDIX A - EVALUATION OF DOD HIGHER LAYER PROTOCOLS

A.1 DOD and OSI Protocol Reference Models

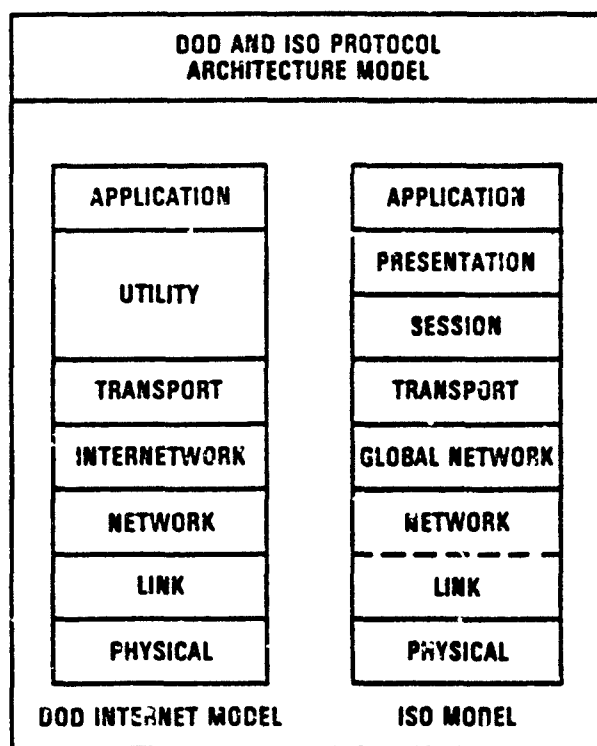
A layered protocol architecture provides a hierarchy of control by functionally decomposing overall network communication objectives into strata. Each stratum, or protocol layer, is supposed to perform a particular function. Starting at the lowest layer, the services of each succeeding layer are available to the layers above and are built upon the layers beneath. The functionality of a layer in no way implies an implementation scheme. Layered architecture allows the protocol designer the freedom to implement the function according to the particular environment and requirements.

Figure A.1 is a simplistic representation of the DOD and OSI layered protocol architectural models. An immediate observation is that the functional decomposition of both is the same through the transport layer. It is only at the higher layers that differences appear.

The reason for this could be that the functional basis for the decomposition changes at this point. Through the transport layer all layer objectives are concerned solely with the transportation of data. A protocol above this layer need not be concerned with physical transmission, routing, involvement or existence of intervening nodes, or errors detected and recoveries made. Instead, above the transport layer, the objective is the accomplishment of a particular task. Both the OSI and DOD models concern themselves with this overall goal, namely, the ability to achieve a common (distributed) task.

Recognition of a functional transition above the transport layer is important. It forces a change of perspective. The higher layer protocols look downward to the transport and lower layers only for the data transmission services necessary for the accomplishment of their distributed tasks. At these higher layers there is less of a stratification of functions. Instead, (and particularly as viewed in the DOD model) there appear to be groupings of decidedly different functions within the same layer. This is because of the common need for data transmission services but the differing purposes of tasks.

An example of this type of functional transition can be found in our postal system. Here the mailbox acts as the interface between the data transportation layers and the task oriented layers. Contents of mailed letters can be viewed as task data. A person mailing a letter is not concerned with the method of transportation, be it truck, train, or plane. The concern is only that



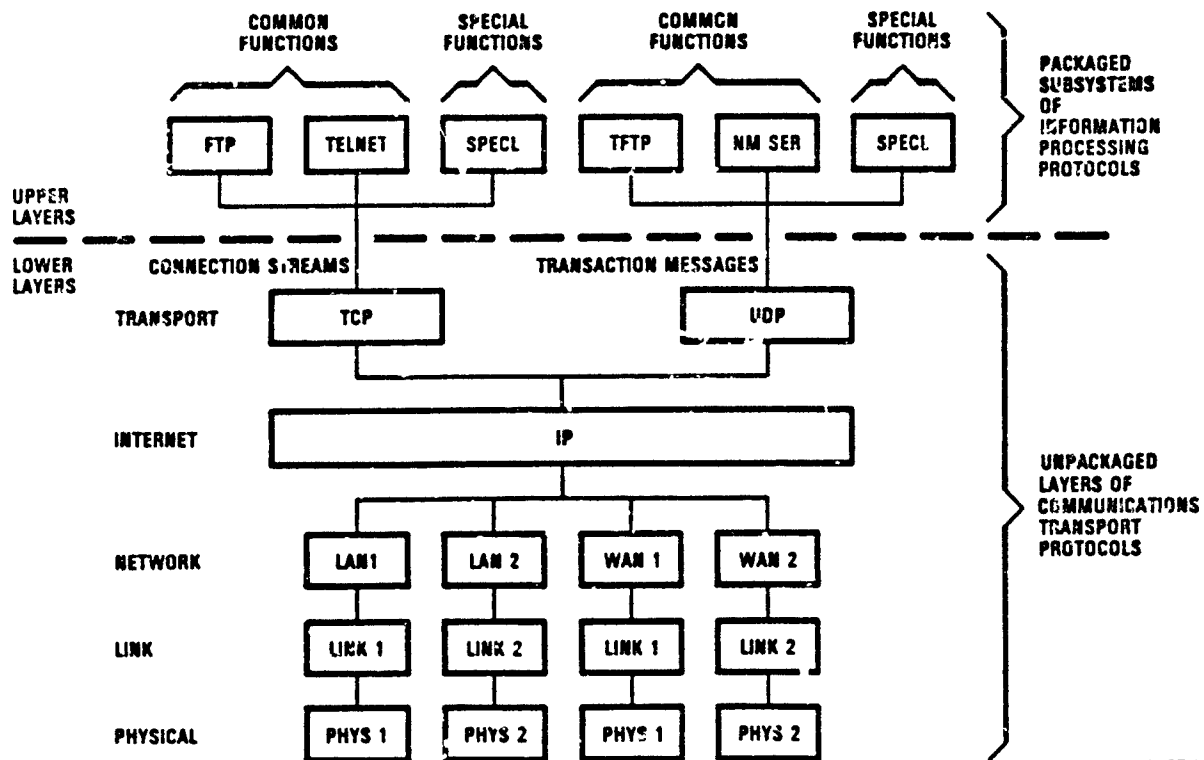
13277-23

Figure A-1. DOD and ISO OSI Protocol Reference Models

the letter arrive at its destination in some reasonable length of time. This is the service upon which the mailer depends. In turn, the postal department has no knowledge of the contents of letters ("tasks"), be they bills, invitations or correspondence.

A.2 Recommended Approach

Figure A.2 is a representation of the recommended approach to viewing protocols at the higher layers. It is based on the contention that a different functional perspective is required above the transport layer. Viewing high layer protocols in this way more closely parallels the real world of distributed information processing tasks. Rather than having each layer above the transport layer represent a single function, instead these layers represent families of protocols. These protocol families represent requirements for specific task operations. Let the application itself (or the user) decide what class of underlying data transport service is necessary, the basic choice being between connection or connectionless support. This approach has been the basis of the IBM SNA Logical Unit (LU) architectural element. Each LU is a tailored grouping of protocol subsets from its Layers 4-7.



13457-51

Figure A.2. Recommended DOD Layered Architecture Model

APPENDIX B

SOME IMPROVEMENTS TO THE DOD HIGHER LAYER PROTOCOLS

APPENDIX B - SOME IMPROVEMENTS TO THE DOD HIGHER LAYER PROTOCOLS

B.1 TELNET - Virtual Terminal Protocol

TELNET is a bidirectional, byte oriented communications facility for terminal-to-process, terminal-to-terminal, and process-to-process communications. It is based upon a scroll mode ASCII TTY. The protocol defines options which are negotiated between parties based upon the characteristics of local devices mapped to a canonical network virtual terminal.

As the protocol now stands it does not provide the capabilities for proper utilization of bit-mapped displays. There are currently many such devices in service. The total facilities of these high resolution displays, such as windowing or font selection, are not served by TELNET.

One approach, which was taken by the Terminal-to-Host protocol, THP, is to define classes of terminals. THP defined four classes of terminals: 1) basic, 2) scroll, 3) paged and 4) forms. By dividing terminals into classes a virtual terminal mapping can be made which provides a better balance between being overly restrictive and overly inclusive. A class could be established for the latest generation of terminals and could provide a mechanism to allow for future technological growth. The problem is a very complex one but one whose resolution is called for as soon as possible.

B.2 File Transfer Protocol - FTP

The objective of the File Transfer Protocol, FTP, is to promote file sharing and encourage the use of remote computers while shielding the user from variations in file storage systems. The protocol is based upon a TELNET connection.

The TELNET connection is used to specify the parameters for the data connection (such as data port, transfer mode, representation type, and file structure) and to specify file system operations (such as store, retrieve, append, delete). The FTP transmission modes are stream, block, and compressed.

The FTP is limited in its ability to allow for the variations which exist between processors in number representation and word length. This is especially noticeable in floating point number representation. It is the intent of the FTP that data transformations beyond those limited ones which are provided be performed directly by the user. As such FTP does not include many built in facilities for file transfers between diverse equipments.

Some file transfer problems are specialized in nature, such as the high speed movement of extremely large volumes of data. These problems might be best addressed by defining new protocols for unique file transfer applications. Aside from specialized file transfer protocols, more thought needs to be directed towards a file transfer protocol which does not require so many user supplied functions. For instance, it may be possible to include type information within the block transfer mode header.

B.3 Internet Name Server

This protocol, given a name string, returns the complete 32-bit internet address as used by the Internet Protocol (IP). The name string is divided into a network portion and a host name portion. The protocol works adequately when there are relatively few networks and names. As more networks become interconnected a centralized name server becomes less desirable and a truly distributed name server becomes necessary. Also undesirable is the fact that a centralized name server creates a single failure point.

A universal network addressing scheme, although for the most part considered highly unlikely to occur, would be extremely beneficial. The telephone industry in this country adopted a universal numbering plan (area code - exchange - subscriber #) whose implementation initially caused quite a commotion. Now we cannot imagine what telephone service would be like if this plan had not been adopted.

A distributed name server protocol should also include facilities for adding, deleting or changing internet addresses of networks and nodes. Also not to be forgotten is the addressing problem created by mobile hosts which can move from network to network.

APPENDIX C

TRANSMISSION CONTROL PROTOCOL (TCP) USAGE FOR LANS

APPENDIX C - TRANSMISSION CONTROL PROTOCOL (TCP) USAGE FOR LANS

C.1 A Case Against TCP

Few people would question the fact that the DOD's Transmission Control Protocol (TCP) is the most robust and thoroughly tested transport layer protocol currently being used across networks. In an internetwork environment where end-to-end reliability and absolute assurance of data delivery is essential a protocol such as TCP must be used. There remains, however, the question of TCP's suitability for intra LAN traffic.

The vast majority of traffic within a LAN is destined for another node (or nodes) within that LAN. That is what a LAN is all about, the sharing of resources and data in order to accomplish related objectives (distributed tasks). Aside from these related objectives LANs may be distinguished from one another according to topology, access scheme, data rate, and types of equipments which may be connected to them. The choice of a LAN is based upon how well a combination of these variables, as exemplified in a specific LAN product, fulfill the overall objectives of the user. There is usually some conscious evaluation of how difficult or feasible internetworking would be, but although a LAN may be rejected because of its inability to internetwork it is not primarily chosen for this reason. The primary goal in LAN selection is to determine how well a LAN suits the specific application level requirements of its client users.

This leads to a myriad of LAN products each based upon different concepts and having different capabilities. It is precisely this diversity (which is desirable from the selector's standpoint) that makes the wisdom of the forced inclusion of TCP within all LANs which internetwork or have the potential of internetworking questionable.

For instance, many LANs possess the characteristic of extremely low error rates. For these LANs, is it really good engineering to include a protocol in the suite for intra LAN traffic which numbers and accounts for every byte of data? Is a resequencing function (as required by TCP) necessary when there is a low probability that messages can be received out of order? Would simpler retransmission schemes suffice rather than the complex TCP procedure?

Put simply, when an underlying network service is reliable, as is the case in the majority of LANs, a protocol like TCP is not necessary. To include it in these cases violates a basic tenet of good design. That tenet is: "Do not provide functions which are unnecessary, which repeat functions already performed, or which are superfluous to the accomplishment of the design goal."

Another argument against the inclusion of the TCP for intra LAN traffic is that it can subtract from the performance of LAN special features. These special features, such as broadcasting and multicasting, may have been a criteria for the original selection of the LAN.

Still another argument against TCP in LANs is the amount of resources and computational overhead it uses. The overhead TCP incurs degrades the overall performance of the LAN. This argument has already been demonstrated. However, as processors become faster and as (virtual) memory size increases this argument tends to dissolve.

The conclusion is that TCP is absolutely essential when internetting but it is generally inappropriate for intra LAN traffic.

C.2 Alternatives

Clifford Warner in his article, "Connecting Local Networks to Long Haul Networks: Issues in Protocol Design (4)", has suggested four alternatives. These are:

1. Make all LAN nodes implement and use TCP for all traffic.
2. Implement two different host-to-host protocols, TCP and a separate network protocol.
3. Place TCP at gateway nodes only and provide for protocol translation there.
4. Implement a local host-to-host protocol which maintains TCP features but which is streamlined for LANs.

Alternative 1 has already been discussed and its disadvantages delineated. Alternative 2 may yield good performance but initial development costs would be great. Additionally, each node would become more complex with the inclusion of two transport protocols. This could lead to design complications and maintenance problems. Alternatives 3 and 4 thus remain as possible candidate solutions.

C.3 The LNTCP Approach

Mr. Warner calls the modified version of TCP in alternative 4 a Local Network Transmission Control Protocol (LNTCP). The LNTCP can be thought of as a protocol which performs a subset of the TCP functions. This subset equates to the major TCP features. Thus translation to a fully implemented TCP at the gateway node is simplified.

To facilitate the translation the mapping of the field definitions which are common to both the LNTCP header and the TCP header should be close to identical. Figure C.3 shows the TCP header and the TCP Psuedo header. The fields which are shaded indicate the most likely candidates for commonality. TCP fields such as the window, checksum and some of the control bits represent the more complex procedures of TCP. They would probably not be necessary in a local network protocol. If any special parameters were required by the LNTCP it is possible that they could be defined as an option and be carried along in the TCP header.

The idea of an LNTCP may at first seem to be a reasonable approach because of the translation simplicity. However, upon further examination, the forced conformity of an LNTCP to TCP features may not be at all practical in a real, operational environment. The worst problem with this approach is the tight coupling of the LNTCP to the TCP. It emphasizes the inter-relationship of these two protocols and as such is an undesirable design goal.

C.4 The Black Box Approach

This leaves alternative 3. In this approach TCP is implemented only at gateway nodes. It is here, and only here, where the execution of a fully implemented TCP takes place. Each internetwork message is encapsulated with a TCP header. (This, in turn, is encapsulated with an IP header, an outbound local network header and a data link header prior to transmission.) The only expectation of the gateway TCP is that a predefined set of interfaces be provided. These interfaces (unlike those in the LNTCP) are not dependent upon any specific underlying protocol.

It can be thought of as a "black box" approach. Except for the gateways the LAN has no knowledge or understanding of TCP. From a system engineering standpoint this is both conceptually and practically desirable. Drawbacks of this alternative are the increased complexity of the gateway node and the possibility of traffic congestion.

When discussing this solution there is frequently a failure to mention what interfaces are required between the LAN local nodes and the LAN gateway node. These interfaces must include all the information necessary for the creation of the TCP header and psuedo header which is not an inherent part of the protocol mechanism itself. Since TCP is a connection oriented protocol, interfaces must also be designed to permit connection establishment and termination. These data items include: source and destination port numbers, the

SOURCE PORT								DESTINATION PORT							
SEQUENCE NUMBER															
ACKNOWLEDGEMENT NUMBER															
DATA OFF SET	RE- SERVED	U R G	A C K	P S H	R S T	S Y N	F I N	WINDOW							
CHECKSUM								URGENT POINTER							
OPTIONS										PADDING					

DATA

I

V

SOURCE ADDRESS															
DESTINATION ADDRESS															
ZERO				PTCL				TCP LENGTH							

13457-52

Figure C-3. TCP Header Comparisons

32-bit internet source and destination addresses (or the information necessary to create these addresses), specification on a per message basis of control bits for urgent and push, a means of returning to the source node the local connection name, and a means of forwarding messages received at the gateway to the local source node.

Two possible ways to provide the necessary interfaces come to mind. In either method each local node would have to perform a test on each outgoing message to determine if its final destination were outside the net. This is necessary because the source node must include enough information for the gateway node to establish whether or not the received message is intended for itself or is to be forwarded outside the net. The source node could then include all the necessary TCP information within its messages. Alternatively, the gateway, upon recognizing a message destined for outside the network, could solicit the source node for any required parameters.

C.5 Conclusions

To unconditionally require the implementation of TCP within all LANs seems inappropriate from both an operational and functional standpoint. The best solution appears to lie in a "TCP at the gateway" approach with no special underlying LAN protocol required.

APPENDIX D
PROTOCOLS IDENTIFIED FOR THE GENERIC NETWORK
OPERATING SYSTEM (GNOS) REFERENCE MODEL

APPENDIX D
PROTOCOLS IDENTIFIED FOR THE GENERIC NETWORK
OPERATING SYSTEM (GNOS) REFERENCE MODEL

D.1 Introduction

The Generic Network Operating System (GNOS) provides a reference model for viewing, understanding and developing protocols needed to build distributed processing systems for U³I. This appendix identifies protocols required for GNOS. Figure D.1 is an architectural illustration of GNOS.

D.1.1 Overview

Message-based Transaction and File Transfer protocols, in support of software program functions, comprise the Generic Network-wide Operating System (GNOS). Software program functions perform distributed Resource (Object) Management operations in support of the distributed applications (policy setting/execution functions reside here) by way of Resource Manager protocols and assessing Networking Utility services.

A networking protocol suite is comprised of three service regions:

- Networking-wide Utilities (canonical form of high level services for accessing remote resources)
- Host to Host/Internet data transport services
- Local/Wide Area Network transmission services

Networking-wide utilities provide generic services to the Resource (Object) Manager entities to enable remote resource access to Files, Terminals, Data, Jobs, Messages, etc. The combination of Resource (Object) Managers and Networking-wide Utilities functions comprises the General distributed Network Operating System (GNOS); where a local resource is employed. The Resource Manager accesses it using its Constituent Operating System (COS).

D.2 GNOS Protocols

Protocols are needed at several levels to enable distributed Object (Resource) Managers to cooperate autonomously. The levels identified consist of the following:

- User to OSCRL*
- OSCRL to Resource Managers
- Resource Managers to Resource Managers
- Networking Utilities to Networking Utilities

* OSCRL - Operating System Command and Response Language (GNOS)

- Host to Host interprocess communications
 - Transport
 - Internetwork
- Subnet Transmission
 - Local Area Network
 - Wide Area Network

D.2.1 Process (Object) to Process (Object) Protocols

Asynchronous handling of simultaneous process to process transaction messages is required. Stream interprocess communications is required between cooperating processes. This has a uni-directional data channel session type between two objects. One is a data source, the other a data sink and connects processes with files, devices, and other processes.

Transactions will be the primary method for objects to communicate through exchanging messages; however, a connection-oriented service will be required for transferring files.

Message exchange characteristics:

- Intra-host (local only)
- Inter-host (LAN only, LAN/Internet/LAN)

Message types:

- Small, minimal effort (datagram service)
- Small, reliable (virtual circuit service)
- Large, reliable (virtual circuit service)

Predominant form of messages exchanged will be for control. These request operations to be performed on objects (local/remote), with replies generated by performed operations, plus exception notices and messages to coordinate the distributed Object Managers.

D.2.2 Resource (Object) Manager to Resource (Object) Manager Protocols

The following Resource (Object) Managers require peer protocols:

- Program
- Terminal Manager
- Authorization/Access Control/Security
- Catalog
- Device I/O
- Network Management
- Directory
- OSCRL
- File

- Monitoring and Control
- Data Base
- Host

Object Manager to Object Manager communications employs a high level form of protocol. It is asynchronous and involves handling interleaved messages from possibly several processes. Messages are received, requests are originated to satisfy the client requests, and a reply message sent to the original message. In the case of failure, the Object Manager assures the client that either all changes requested will take place, or none will for the atomic transaction performed. Standardized Object Manager to Object Manager message types and structuring is required of message formats employed.

Resource Managers make a set of resources available to users (programs), such as processor cycles, main storage, files on disk or tape, such I/O devices as keyboards or displays, and such abstract resources as sessions, queues, or data base records. Resource Managers allocate resources to users as its central function in response to a user request. This includes access scheduling, coordination, resource allocation deadlock detection, resource change commitment control, resource access security and resource formatting services. Resource Managers employ resource coordination peer protocols and access network services by way of interfacing to the Networking-wide Utility protocol services. Program to Program protocols exchanged between Resource Manager/Network Utility entities make up the distributed network operating system.

Protocols support cooperation among the distributed Resource Managers to perform the following activities:

- Interprocess communications
- Data representation
- Data storage (media, file and data base)
- Process management
- Resource management
- Integrity and security
- Program support

D.2.3 Networking-Wide Utility* Protocols

- Network Management
- Virtual Terminal
- File Access, Transfer, Management

* Supported by its type Presentation and Session protocols.

- Job Transfer/Manipulation
- Message Handling
- Document Interchange
- Name Server
- Data Base Access/Management

D.2.4 Interprocess Communications

Interprocess communications facility exhibits a loosely coupled message passing characteristic (not memory sharing)

- Transaction (connection-less) to be predominant form of exchange
- Stream (connection-oriented) will also be required

D.2.5 Gateways Elements are Required for Interoperability

Intrasystem and intersystem connectivity functions are required to be performed.

APPENDIX 2
NETWORKING AND SYSTEM RESOURCE MANAGEMENT
IDENTIFIED PROTOCOLS

APPENDIX E
NETWORKING AND SYSTEM RESOURCE MANAGEMENT
IDENTIFIED PROTOCOLS

E.1 Introduction

An important aspect of Open Systems Interconnection (OSI) is the organization of the distributed processing activity and the resources required for its successful prosecution. The work on Application and Systems Management, and Job Transfer and Manipulation deals with the specifics of distributed processing activities. This appendix identifies management protocols required and some characterizations of their functionality.

E.2 OSI Systems Resource Management

This, in conjunction with Job Transfer and Manipulation, deals with the organization of Distributed Processing activity and the resources required for its successful prosecution.

Examples of OSI Systems Management include:

- Regular management activities for:
 - Resource allocation/deallocation
 - Access control
 - Process activation/deactivation
 - Accounting
- Change management
 - Reconfiguration
 - Name handling
- Security management
 - Authentication
 - Encryption
- Integrity management
 - Including commitment control
- Error reporting, recovering, and journaling

E.2.1 OSI Management - Deals with Managing

Application-Process-Group, Systems and Layers as follows:

1. Application-Process-Group Resource Management

- Control and monitoring of an Application-Process-Group's activities
- Objects being managed are conceptual views of: Files, Data Bases, Job Processing Resources, etc.
- Functions include concurrency control, recovery, accounting

2. Systems Management

- Control and monitoring of the NOS networking utilities resources (OSI resources)
- Functions include resource activation/deactivation, allocation and deallocation, and reconfiguration

3. Layer Management

- The control and monitoring of the communication resource

E.2.2 Applications-Process-Group

This is a collection of information processing activities (application processes) which collaborate to meet a specific information processing need. Within the OSI layered architecture, this is manifested as a related group of communication activities between application entities.

E.2.3 An OSI Resource

This is a conceptual view of a class of actual resources that may be used by the Application-Process-Group in meeting the user's information processing requirement. Instances of OSI resources are termed management objects.

E.2.4 OSI Management Protocols

These constitute the rules by which management information will be exchanged among the end and open system resource managers. The OSI management framework has been divided into two key concepts; the Applications-Process-Group and the OSI resource.

E.2.5 Some Specific Management Standards Topics (Future Work Required)

- Directory information
 - Names, addresses and attributes of OSI resources
- Accounting management
 - Information on charges for use of resources
- Authorization management
 - Information about management authorization
- Commitment, concurrency, recovery
 - Ensure integrity of information processing against system malfunction
- Control of application-process-groups
 - Monitor and control activity for enrollment/de-enrollment, activation/deactivation, initiation/termination of application-process-groups

E.3 Job Transfer and Manipulation (JTM)

JTM services and protocols facilitate job processing in distributed systems, relating not only to the movement of job-related data between open systems, but also to the movement of information to monitor and control job processing activity. A job in OSI is described in terms of the type of work to be carried out.

E.3.1 JTM System

Is subdivided into four functional components:

- Job Submission System (issues demand for job)
- Job Processing System (does the job)
- Job Monitoring System (advises on job progress)
- Job Manipulation Submission System (controls the JTM activity)

E.3.1.1 JTM Entities

Communicate through the transfer of work specifications, which are structured data objects that define all the work which the recipient is required to perform.

E.3.1.2 JTM Work Specification

Consists of two parts; a control part and a documents part. The control part contains work control information and are subject to JTM protocol standards. The documents part is transparent to JTM and conveys recipient specific information. JTM acts on the information in the control part but passes the documents to the local environment (end system).

E.3.1.3 Management of Application-Process-Groups

Both OSI Management and Job Transfer and Manipulation have features which are common and relate to the management of APG's. On the one hand, OSI Management deals with the whole of distributed management, whereas JTM is dealing with the specific.

E.4 Open Systems Management Issues

E.4.1 OSI Management Control of Application-Process-Groups

This deals with standardized access to, and use of, the information processing resource and the protocols to initiate and monitor the activities of an APG as an instance of an OSI resource to be managed.

E.4.2 Job Transfer and Manipulation

This environment presupposes existing job processors to which work specifications are submitted for a specific piece of job processing. Separate subjobs are only loosely related in time, and the structure of the distributed job does not rely upon the prior commitment of resources to accomplish each subjob.

E.4.3 General Distributed Processing Management

This activity requires considerably more generality and functionality than currently possessed by the JTM service. A uniform Operating System Command and Response Language (OSCRL) has been identified for the general model of system control and monitoring functions in a generalized distributed processing environment. This also has ramifications on work under way for programming languages aimed at standardization of the user's language interface with a distributed processing activity. It is essential that the distributed processing models adopted by the OSI, OSCRL and Programming Languages groups correspond with each other. Failure to do so would result in incompatible multiple standards and their accompanying costs.

E.4.4 Implications of Open System Interconnection

Where current work in OSI has dealt with communications protocols for seven layers of services, the implications of OSI extend far beyond the instance of just open communication. Experts are now, in OSI, viewing how the OSI layered approach may be viewed in relation to the full scope of distributed information processing and not just communications. It has been proposed in ISO to restructure the current work on Programming Languages, OSCRL and communications into a Global OSI project spanning Distributed Information Processing.

APPENDIX F
REMOTE DATA BASE ACCESS PROTOCOL

APPENDIX F

REMOTE DATA BASE ACCESS PROTOCOL

F.1 Introduction

A set of protocols is required within a distributed processing system to enable remote access to data located at another location in the network. During the study, two works reviewed were found to provide a framework for structuring the development of suitable protocols; the following briefly discusses these.

F.2 Interconnecting Heterogeneous Data Base Management System

Reference [98] presents a discussion on the analysis of the existing approaches to interconnecting heterogeneous Data Base Management Systems (DBMS) and reviews alternatives from four experimental projects. An architectural model presented is shown in Figure F.2.

The functional layering of the network of heterogeneous DBMS's, which refers only to the application layer of the ISO reference model described by the International Standards Organization, consists of three sublayers: 1) the Global Data Manager, or GDM, is the top-most sublayer that provides services directly to the end user; 2) the Distributed Transaction Manager, or DTM, is the middle sublayer that supports the services of the GDM and requires the services of the Structured-Data Transfer Protocols; and 3) the Structured Data Transfer Protocols, or SDTP, is the lower sublayer that supports the services of the DTM and requires the services of the Data Presentation Protocol, or DPP, and of other application layers, such as those of the File Transfer Protocol, or FTP.

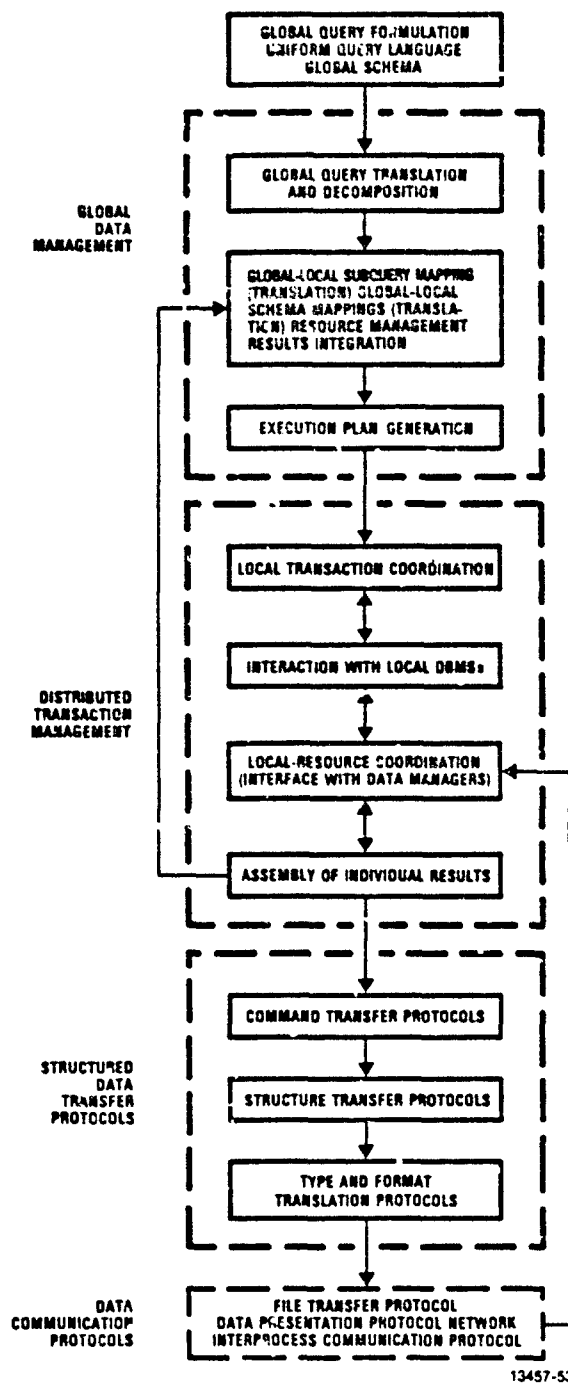
F.2.1 Global Data Manager

The Global Data Manager (GDM) of a distributed DBMS performs both the mapping between the global unified view of the data and the local DBMS's, and all relevant I/O operations. The following five functions are included in the GDM:

1. Global data model analysis
2. Query decomposition
3. Query translation
4. Executing plan generation
5. Results integration

F.2.2 Distributed Transaction Manager

The Distributed Transaction Manager (DTM) is responsible for controlling the execution of distributed transactions in integrated distributed DBMS's; those transactions will reference data at more than one site in a network. A primary purpose of a transaction manager, whether it is distributed or centralized, is to help ensure the consistency of the data base. To do this, the



13457-53

Figure F.2. Architectural Protocol Model for Heterogeneous DBMS

DTM assumes that whenever a transaction runs in isolation and completes its task, it preserves the consistency of the data base. Concurrency control schedules the subtransactions of concurrently operating transactions. Recovery control provides data protection and recovery procedures for all transactions.

F.2.3 Structured Data Transfer Protocols

The Structured Data Transfer Protocols (SDTP's) are application-level (Layer 7) protocols required for the interconnection of remote heterogeneous DBMS's. The SDTP's are used by protocols and programs at higher levels which implement the interconnection of the remote DBMS's. These protocols and programs belong to the GDM and to the DTM. The SDTP's themselves use the data communication protocols. The areas covered by the SDTP's are the following:

1. Command Transfer Protocols
2. Structure Transfer Protocols
3. Type and Format Translation Extension to a Data Presentations Protocol

F.2.3.1 Command Transfer Protocols

The Command Transfer Protocols deal with the transfer of commands to remote DBMS's. These protocols are particularly useful in cases where some of the transaction and data management protocols are implemented as remotely located servers, rather than as application-protocol layers. As with other SDTP's, the Command Transfer Protocols use a canonical command format for command transfer across the network. Higher levels, such as the GDM, will perform the translation between local command formats and the canonical form.

F.2.3.2 Structure Transfer Protocols

The most important SDTP's are the Structure Transfer Protocols. A structure can be a Pascal or a Cobol record, a PL/I, or a C structure. The Structure Transfer Protocols include two basic kinds of protocols: 1) protocols for structures with pointers, and 2) protocols for pointer-free structures. The STP is a direct user of the File Transfer Protocol and Data Presentation Protocol layers since files are most likely to be the representation type for most structures with or without pointers.

F.2.3.3 Type and Format Translation Protocols

Type and Format Translation Protocols are necessary for types of objects not included at the DPP level. The Structured Transfer Protocols require that the DPP level must be extensible, that is, the DPP must be prepared to handle a large variety of types in addition to four currently specified.

Reference [58] discusses functions of the SDTP protocols in more depth.

F.3 Remote Data Base Access Protocol of ECMA

F.3.1 General

The European Computer Manufacturers' Association (ECMA) Standards Organization has developed a draft Remote Data Base Access Protocol [99]. This was deemed by the study to represent a major contribution. The principal stated characteristics of this protocol are that:

1. It offers efficient remote access to data base
2. It supports a distributed data base

It provides a general framework for remote access to data bases of many types, with specific encodings for access using the Relational Model.

The protocol standard is one of a set of standards for Open Systems Interconnection. It is intended to facilitate homogeneous interworking between heterogeneous information processing systems. It is consistent with emerging data base standards being created by ANSI X3H2, and capable of modular extension to cover future needs and to exploit future developments in technology.

The Standard for the Remote Data Access Protocol:

- Defines a data base model
- Defines the operations on the data base model as abstract interactions between two users of the communications service, one of which is acting on behalf of an application program while the other is interfacing to a process that controls data transfers to and from the data base
- Defines the protocol to support the above service and its mapping to the underlying presentation service
- Specifies the requirements for conformance with this protocol

The protocol is targeted at data base systems that support the relational model. However, it is envisaged that other data base systems will support relational interfaces and that subsets or supersets of the protocol may be used with other data base systems.

This protocol is for the Application Layer of Open Systems Interconnection.

F.3.2 Data Base

This section describes the characteristics of a data base that are assumed by the model and identifies specifically those that are visible to remote users of the data base. The DBMS consists of all the compilers, utilities and data base access software necessary to support the creation, management and use of the data base. The term Data Base Control System (DBCS) specifically refers to

run-time component providing data base access and manipulation facilities to application programs.

F.3.2.1 The Data Base and Distributed Data Base Models

F.3.2.1.1 General Principles

A data base is a coordinated body of data managed by a software entity termed a Data Base Management System (DBMS). The DBMS maintains the data base in permanent storage. It controls and facilitates the storage and retrieval of data in the data base.

A logical structure of the data base is defined by a data base description known as a Schema. This defines the names and characteristics of all the data elements that may be stored, the interrelationships among data elements, and the constraints on the values they may take. Details of the mapping of data to storage and performance requirements are generally defined in the Internal Schema or Storage Schema so that the logical capabilities that are visible to programmers can be separated from the performance concerns which are the responsibility of a data administrator.

Access to the data base by application programs may be via a procedural interface or by extensions to a Programming Language which we can imagine being compiled into code that invokes the same procedural interface. Since different application programs require access to different subsets of the data base the data available at the procedural interface for any single connection is defined in a separate data definition: the Subschema or External Schema. The procedural interface specification defines the data manipulation functions available and their effects on the data base.

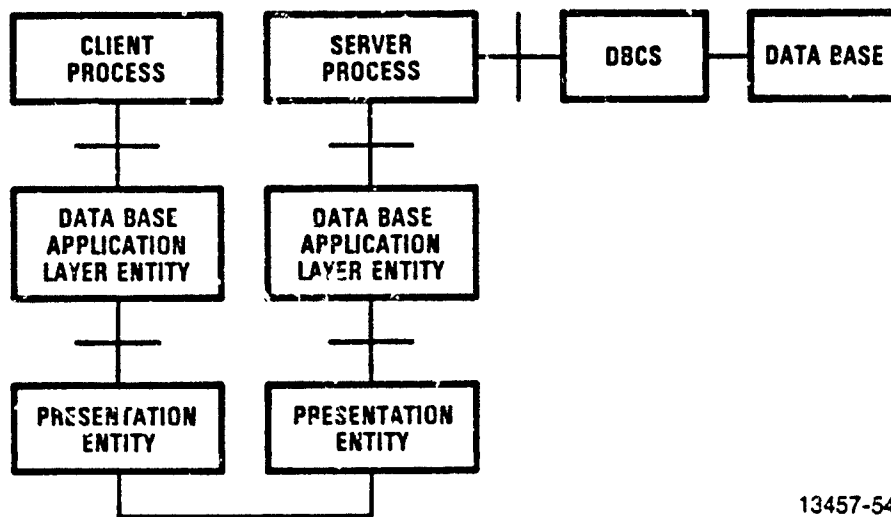
Specifications of the data structure descriptions and the data manipulation functions have been developed by ANSI X3H2. The semantics of the functions described in this specification and their effects upon the data base are defined in that document.

The current ANSI RDL does not include a Subschema. However, the subschema required by this protocol may be identified to the RDL schema so there is no conflict. Also, the subschema invoked in RDA may be a relational data base description which an implementor has provided for nonrelational data bases.

F.3.2.1.2 Remote Access to a Data Base

Figure F.3.2.1.2 shows the structure for remote data base access.

The client process is an application program or Query Language Processor that has a data processing job to do. The Application Layer entities are software components that handle the communications on behalf of the client



13457-54

Figure F.3.2.1.2. Structure of Remote Data Access

process and the server process. At the data base end, the server process is translating the protocol messages into Data Manipulation Procedural Interface calls and parameters and transmitting the results back using the service primitives.

The diagram does not show the structure of the client process. However, it is envisaged that the service interface on the client side will generally be driven by a component of a DBMS (or distributed DBMS) so that the user interface for remote access to data is not unnecessarily different from the interface used when data is local.

F.3.2.1.3 A Distributed Data Base or Multi-Data Base System

A distributed data base is a coordinated body of data that is partitioned into separated data bases, each managed by its own DBMS. Coordination across the components is managed by a distributed DBMS (DDBMS) which is itself distributed.

It is possible for a client process (or application program) to access the distributed data base without being knowledgeable about the location of the

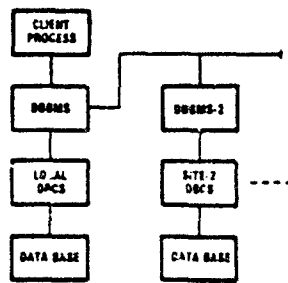
data elements. Figure F.3.2.1.3-1 shows a client process calling a DDBMS component in its own end system which communicates with peer entities (DDBMS components) in other end systems via the DDBMS connections.

Each connection between DDBMS components is an OSI Application layer connection. The transport connections will handle many transactions simultaneously, but the protocol defined in this document is concerned with the interaction between the site with the client process and one other site. Figure F.3.2.1.3-2 shows the structure of a single connection.

F.3.3 Protocol Characteristics Summary

The Remote Data Base Access Protocol document addresses the following characteristics:

1. Service Description
 - a. Roles of Partners
 - b. Dynamic Structuring of an RDA Connection
 - c. Connection Services
 - d. Subschema Management Services
 - e. Data Manipulation Definition Services
 - f. Transactions
 - g. Data Manipulation Functions
 - h. Bulk Data Transfer
2. Protocol Specification
 - a. Connection Management
 - b. Transaction Management
 - c. Grouping of Statements
 - d. Bulk Data Transfer
 - e. RDL Statements and Macros
3. Conformance



12087-00

Figure F.3.2.1.3-1. Distributed Data Base Schematic

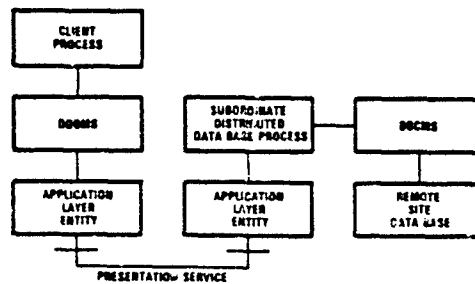


Figure F.3.2.1.3-2. Distributed Data Base Single Connection

APPENDIX G

GENERIC GATEWAYS FOR LAN INTEROPERABILITY

APPENDIX G

GENERIC GATEWAYS FOR LAN INTEROPERABILITY

G.1 Introduction

Gateways are functions which join two networks together. A gateway may look to each network like one of its own nodes; on the other hand, a special purpose function node may be specified. A gateway comprises the collection of hardware and software required to effect the interconnection of two or more networks enabling the passage of user data from one to the other.

G.2 Generic Gateway Family

Gateways are utilized to achieve internetworking. Here, internetworking refers generally to the interconnection of networks, whether similar or dissimilar. Where internetworking is accomplished at a given layer of the OSI/RM, generally then requires any two applications Layer 7 users who actually wish to communicate and interoperate must choose and employ common protocols for the layers above where internetworking is performed and up through Layer 7. If this is not done, then one must find and employ a resource to perform protocol conversion functions between the two noncompatible protocol suites.

The solution to the interconnection of different types of local networks, and the connection of local networks to long-haul networks, is the use of high-performance networking gateways and bridges. A gateway server traditionally operates between two similar or dissimilar networks and can communicate with the system elements or nodes on each of the adjacent networks. A gateway also performs routing or protocol translation functions that allow some level of internetwork communication among system elements. Gateways are naturally balanced systems, each implementing two half-gateway functions, plus a common relay operation to join the two halves together.

A family of generic gateways is needed to span the full range of the OSI/RM's seven layers. Table G.2 lists the generic gateway set. Each is discussed in the following paragraphs.

G.2.1 Amplifier Gateway (Layer 0)

This type of gateway is utilized to extend the length of a media segment employed in a LAN. It is a linear device which amplifies the received signals and may compensate for impairments experienced by the signal received prior to retransmission. This operates on the composite signalling bit symbols contained within the bandwidth capacity of the two media being interconnected. This can be either unidirectional or bidirectional. Higher layer protocols are transparent to the effects of this gateway.

Table G.2. Generic Gateway Types

<u>Gateway Type</u>	<u>OSI/RM Layer(s)</u>
● Open Systems Interconnection	7-0
● Protocol Conversion	7-4
● Internet	3-0
● Network Relay	3-0
● Bridge	2-0
● Repeater	1-0
● Channel Translator	0
● Amplifier	0

G.2.2 Channel Translator (Layer 0)

This type of gateway is utilized to connect the frequency of one LAN channel to a different frequency of another LAN channel. This is a linear device which performs a frequency shifting (up or down) and may amplify and compensate the signal present on the medium employed. This is analogous to the Amplifier Gateway in that it extends the length of the LAN. This gateway may be statically configured or dynamically switchable to different channel frequencies. Except for address assignments associated with the channel identifications, all higher layer protocols are transparent to the effects of this gateway.

G.2.3 Repeater Gateway (Layer 1-0)

This type of gateway is utilized to extend the length of a LAN's physical topology. It is a nonlinear digital device and converts the received digital bit stream into a reshaped and regenerated outgoing bit stream. The underlying media and signalling methods may be the same or different. This may be employed to interconnect two in-house metallic cable based LAN media by way of a fiber-optic medium running outside and between the locations or any two remote locations by way of a leased point-to-point circuit employing modems. Appropriate control signals may need to be extended between the two sites. Higher layer protocols are transparent to the effects of this gateway and it differs from the Bridge Gateway in that the Repeater Gateway stores and forwards bits, not packet frames or packet messages.

G.2.4 Bridge Gateway (Layers 2-0)

This type of gateway is a packet frame store and forward device which is utilized to join LAN segments together containing either the same or different MAC protocols (i.e., CSMA/CD to/from Token Bus Medium Access Control). This contains two half-gateway MAC's plus a relay function which joins the two halves

together. All of the normal MAC receive de-encapsulation functions must be performed (i.e., framing, error detection, address recognition, and the particular protocol's control functions). Then the relay function may perform a filtering function to only pass on frames containing addressed data destined for a station on or beyond the next LAN's MAC segment being interconnected. Selected addressed frames are then encapsulated with the structure of the new LAN's MAC protocol structure and when appropriate is transmitted.

Since all of the IEEE 802 MAC protocols operate in the connectionless service mode, each MAC frame processed is treated as a separate entity. Whether the next sublayer (LLC) is providing connectionless or connection-oriented service, the Bridge Gateway appears transparent to the LLC packet frames. The LLC originating and destination stations manage their respective state(less) and data management functions on a LAN end-to-end basis, just as if the Bridge Gateway was not present. For those cases where the two MAC protocols joined at the gateway do not provide equivalent services, then either functionality must be added to the protocol of lesser equivalence or the most common set of equivalent services is used to form the service basis for interoperability (i.e., the higher grade of service is subsetted down to the lower's equivalence).

This type of gateway is used to extend the range and overall capacity of the original LAN(s) employed. Higher layer protocols can be made transparent to the effects of this gateway when a proper set of equivalent services is achieved and used in the interconnection mode.

A higher layer version can also be employed by a Bridge Gateway. Here, the gateway terminates the MAC and LLC services of one LAN, and performs a pure relaying operation to extend this LAN's LLC services onto the next LAN's LLC protocol sublayer. This is a form of cascading like protocol services in a serial string arrangement. With connectionless LLC service, individual LLC frames are passed transparently by the relay to the next LLC protocol entity. However, for connection-oriented LLC service, each LLC logical link terminates at the receive half of the gateway, the information field is then relayed into a link connection of the next LAN's LLC entity, and a new set of station-to-station state variables is established for each succeeding LAN in the cascade. Flow control, sequencing and recovery would span at most a single LAN at a time, since LLC is not a global end-to-end protocol, but rather a single LAN's span. This LLC relay form of Bridge Gateway would avoid the use of a separate Internet Protocol above the LLC sublayer as is employed in the Internet Gateway and can be employed where each underlying MAC protocol entity, media and topology elements differ.

G.2.5 Network Relay Gateway (Layers 3-0)

This type of gateway is intended for use in interconnecting a LAN to a WAN and/or a WAN to another WAN, without having to employ an Internet Protocol sublayer above it (see Internet Gateway below). In the case of interconnecting two X.25 packet switched Public Data networks together, the CCITT X.75/X.121 combination of recommendations perform this Network Relay type of Gateway operation. There, X.75 performs a relaying operation where origin network X.25 virtual circuit connections are mapped onto new X.25 virtual circuit connections of the destination network in accordance with cascading principles. There may be three or more state-dependent virtual circuits plugged into a series arrangement of an end-to-end basis.

A proposal [67] made by the author in 1980 to develop a gateway standard approach for interconnecting IEEE 802 LAN's by way of X.25 WAN's has been translated by H. C. Folts into a pragmatic approach [68] based upon the Network Relay Gateway architecture. In this approach, an X.25 Network layer set of protocol functions is added on top of the IEEE 802 suite of LLC, MAC, protocol layers and relayed onto the X.25 Network layer used to access a wide area network. There, the X.25 Network layers employed in the 802 LAN's as well as the X.25 WAN(s) would each operate in the connection-oriented service mode on an end-to-end cascaded basis (like the way X.25 - X.75 - X.25 networks do in series). Above Layer 3 an appropriate Layer and Transport protocol would be employed.

G.2.6 Internet Gateway (Layers 3-0)

This type of gateway is employed to create a new global network out of the interconnection of otherwise heterogeneous LAN and WAN subnetworks. An internet-working sublayer protocol is incorporated in each source and destination station plus the gateway nodes which join a LAN to WAN. Generally, connectionless operation is employed to do the functions of end-to-end packet routing, switching and fragmentation-reassembly of packets which are too large for a subnet. The Department of Defense Internet Protocol (IP) is the best example of this gateway approach. Within each gateway, a relaying, routing and switching operation is performed on the control information contained in each packet header. The Internet Gateway, employing a connectionless protocol, is a more general and robust solution to internetworking than the above discussed Network Relay Gateway, in that it can accommodate a heterogeneous mix of subnets, each being either connectionless or connection-oriented.

G.2.7 Protocol Conversion Gateway (Layers 7-4)

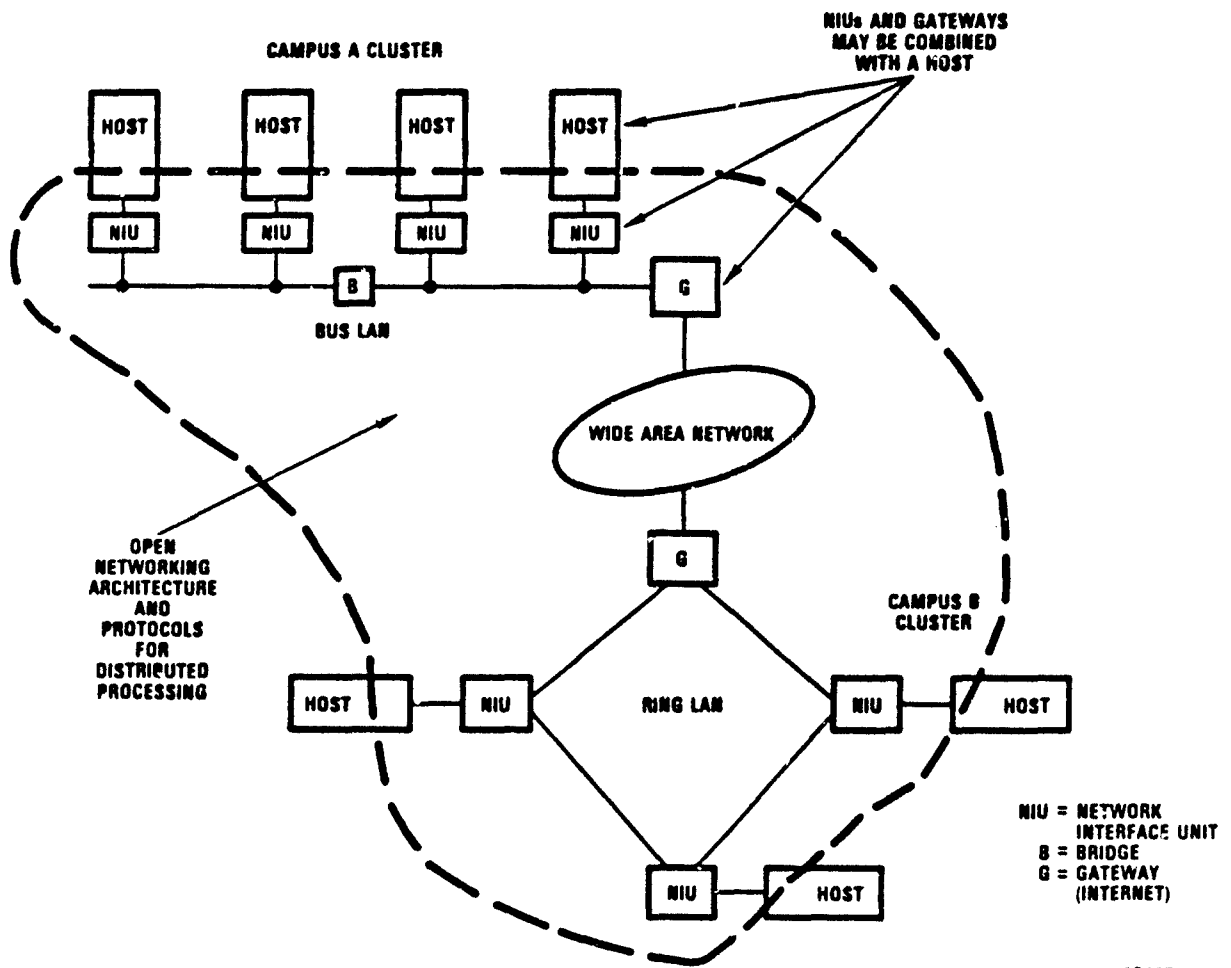
This type of gateway is used to interconnect networks that have different protocol architectures (i.e., IBM's SNA with DEC's DNA). Basically, this might span any range of protocol layers, but it is at the upper layers where the greatest differences are expected to occur from the variations in proprietary solutions being offered in the marketplace. Proprietary vendor protocols have been developed to structure the internal architecture of a system, this being a closed rather than an open form of end system, which the OSI/RM's protocol suite has been defined to interconnect. Therefore, the Protocol Conversion Gateway has to be structured with protocols from the two proprietary networks attempting to be interconnected. Further, a place up in the higher protocol layers must be found where the equivalence of services either exists naturally, or by functional enhancement/de-enhancement can be made to be equivalent in semantics. Thereafter, translation of syntax formatting of the control primitives and information representation has to be performed to map between the two protocols. Most likely, this will need to be performed at or above the Presentation Layer. In addition, mapping may be necessary between the two networks' resource management protocols as well.

The Protocol Conversion Gateway is the least likely one to ever become standardized, because for every combination of vendor network to be interconnected, there is a different design which this type of gateway has. Where implemented, it is desired to do this only once in a gateway node which joins the two networks together. When two protocols are not "connectable," the only possible solution for communicating with the other entity is actually to implement the other entity's protocol.

G.2.8 Open Systems Interconnection Gateway (Layers 7-0)

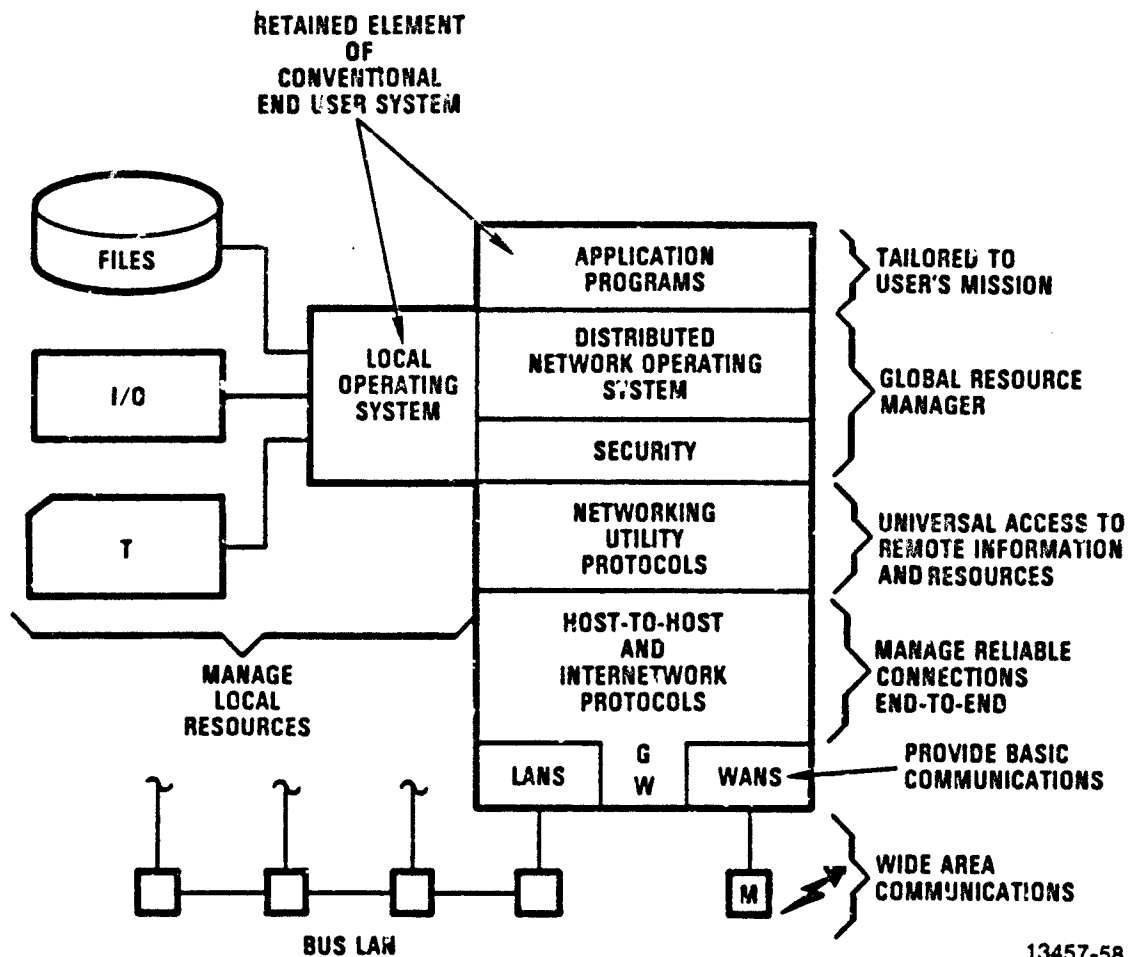
This type of gateway enables the interconnection and interoperability of otherwise closed (proprietary architecture/protocol) end systems by forming an open global system of systems. The OSI/RM provides the international architectural agreement on the structure of protocols and their rules of operation to enable open systems cooperation. An end system which obeys applicable OSI standards in its cooperation with other systems is termed an open system (Figure G.2.8-1 and G.2.8-2).

From a using end system's perspective, the OSI provides an end system to end system set of virtual resource management protocols (for accessing terminals, processes, files, jobs, messages, documents, data, OSI resources) as well as how to communicate data through a concatenation of LAN's and WAN's via



13437-57

Figure G.2.8-1. Unified Networking Approach



13457-58

Figure G.2.8-2. New Distributed Processing Architecture

gateways. The OSI protocols only deal with the exchange of information between end systems and not with the internal functioning of each system.

The application of the OSI/RM with its suite of protocols will foster the use of universally agreed-upon means of permitting communication and cooperation between (or among) heterogeneous systems and products. Existing systems will progressively implement OSI capability in response to user application needs. This will lead to increasing diversity of heterogeneous open systems; heterogeneous because they are built on different architectures; and open because they are capable of cooperating with other systems by implementing the OSI protocols.

G.3 References

The references applicable to the work discussed above are [62-69].

APPENDIX H
MULTIMEDIA LAN (MMLAN) INTERNETWORKING

APPENDIX H

MULTIMEDIA LAN (MMLAN) INTERNETWORKING

H.1 System Concept - MMLAN

This discusses a preliminary system baseline concept for the integration of the FILAN with a possible MMLAN. This is illustrated in Figure H.1, which depicts a hypothetical distributed command center.

Current tactical nodal centers of the Air Force are highly centralized and therefore are vulnerable to enemy attempts at destruction. This threat along with new LAN and distributed system technologies will permit distributing the resources of a centralized nodal center outward among interconnected shelters and vans. The FILAN can provide the communication facilities for use within each shelter while the MMLAN can interconnect together, with multiple serviceable links, the collection of shelters to form an integrated operation center.

In Figure H.1, a gateway node is intended to interconnect individual FILAN Local Subnets to the multiple media links which span between the FILANs. The gateway would function as a Network Access Unit (NAU) on the FILAN to which it belongs as well as be a user of each data link provided by the MMLAN. Functionally, the gateway would incorporate protocols of the physical, data link and internetworking layers associated with FILAN and the DOD Internet Protocol suite on the FILAN side, and individual media, physical and link protocols from the new MMLAN suite. A collection of Network Management functions and protocols would also be a part of each gateway with a close operational tie to each FILAN Network Management node.

Within the MMLAN subsystem's sphere, there would be employed a mixture of terrestrial, airborne and free-space data link transmission facilities. These would be configured to satisfy the specific constraints of the particular deployment in such a way to ensure survivability in the event any data link, FILAN subsystems or both were destroyed. Currently envisioned examples of media types for potential use are as follows:

1. Fiber optic¹
2. Microwave radio¹
3. Spread spectrum data links with a remote pilotless vehicular node (airborne)²
4. Millimeter wave radio²

¹ Denotes near time frame.

² Denotes longer term.



五

5. Infra red radio2
6. Packet radio2
7. Terrestrial1
8. JTIDS2

Of all these media types, the concept of employing a terrestrial satellite (based upon an overhead circling RPV node employing spread spectrum techniques) appears to be well suited to providing full connectivity among distributed nodal shelters. Such a configuration appears to be well suited to the objective of achieving a distributed, survival command center to counter foreseen tactical threats. Harris Corporation has conducted studies in the use of the MICNS RPV data link system to perform such a mission and has documented a systems design approach in a classified report. On the other hand, low cost packet radio appears to be another viable alternative for consideration. The technology has been demonstrated by experiments conducted by the army at Ft. Bragg employing elements of the TCP/IP Internet Protocol suite. Two other media appear to offer substantial benefits in their own right; fiber optic cable and millimeter wave radio. The costs and performance capabilities of these will have to be examined further. In any real system deployment of the MMLAN, no single media would be relied upon exclusively to ensure survivable communications. The system design would be structured to be able to mix and match any of the selected media types then exercise operational systems control to maintain the minimum essential degree of availability specified for the mission.

Harris Corporation is currently under contract [100] studying the application of LAN and multimedia for a MMLAN system definition.

END

FILMED

6-85

DTIC